THE THEORY OF THE EARTH–MOON LOGISTICS REGIME

By Ernesto Gabriel Arias Fernandez

Version 1.1
First edition. Document subject to revisions, corrections, and expansions without modification of its fundamental structure.

Barcelona – 2026

AUTHOR'S NOTE

This work was born from a discomfort. Not with people, nor with institutions, nor with space programs, but with the way the fundamental question is formulated.

For decades, lunar exploration has been conceived in terms of vehicles, missions, and campaigns. However, a permanent human presence is not a mission: it is a ***system****. And a permanent system requires a* ***permanent logistics structure****.*

When one changes the question—when one stops asking "How do we get there?" and starts asking "How do we remain indefinitely on the Moon?"—the answers change. And some of those answers are uncomfortable.

Applying ***physics****,* ***logistics****, and* ***systems theory*** *to publicly known space architectures produces results that do not always align with the language in which those architectures are presented.*

A theory that avoids pointing out what does not work is not a theory: it is an academic exercise. This work does not accuse anyone of bad faith. It is an ***audit of applied physics*** *applied to ambiguous public language. The data are public. The equations are the same as always. The only thing that changes is the* ***conceptual framework****.*

This work is foundational, not definitive. It describes a ***logistics regime*** *for a specific case: the Earth–Moon network under conditions of dominant asymmetry. But the mathematical object that emerges from these pages is more general.* ***Unidirectional Networks*** *are reusable. The* ***NETWORK****—the complete graph of space transportation in the solar system—is the natural generalization.*

The conjecture of oscillation between regimes, the formalization of multiple nodes (Mars, L2, stations in LEO), the integration of asymmetric ISRU, the extension to multi-operator architectures with open standards: all of that remains to be written.

Someone, at some point, will have to do it. Perhaps it will be you.

If this theory is useful to you—whether to build, to analyze, or to demonstrate that it is wrong—then it has already fulfilled its purpose. Because in the end, this work is not mine. It belongs to whoever takes it and carries it further.

PROLOGUE

This entire work was born from a single question:

What logistics structure does a system need in order to sustain a permanent human presence?

The first problem was not finding the right answer. It was finding the right question.

For decades, lunar exploration was studied as a succession of vehicles, missions, campaigns, and programs. But a permanent human presence does not constitute a mission: it constitutes a **system**. And a permanent system requires a **permanent logistics structure**.

The question therefore ceased to be *how to reach the Moon*. The real question became:

What properties must a logistics network possess in order to sustain a human presence beyond Earth indefinitely?

That difference, seemingly small, completely changes the object of study.

It is important to note that the present theory does not propose an Earth–Moon logistics network as an architecture designed by the author. It asserts a different thesis: that such a network is a **necessary consequence** of the physical laws governing space transportation when the objective is to establish a permanent human presence.

In this sense, the theory does not describe an invention, but a **mathematical structure** implicit in the Tsiolkovsky equation, in orbital mechanics, and in the general laws of logistics. The author's work consists of identifying that structure, formulating it through axioms, defining its properties, and deriving its consequences.

The theory does not create this object. It simply attempts to describe it.

If the hypotheses and demonstrations are correct, then under the conditions of **dominant asymmetry** described in this work, any logistics architecture that attempts to sustain a permanent human presence between Earth and the Moon will tend to converge toward a structure with these same fundamental properties, regardless of the vehicles, technologies, or materials employed.

For that reason, the object of study of this work is not the vehicles. Nor the missions. Nor the specific architectures.

The object of study is the **Earth–Moon Logistics Regime**, understood as a mathematical structure whose existence derives from the very laws of space transportation.

This prologue is an invitation: to leave behind the epic and enter the domain of systems. The Moon does not need heroes. It needs infrastructure.

THE THEORY OF THE EARTH–MOON LOGISTICS REGIME

CENTRAL POSTULATE

Establishing a bidirectional logistics system between Earth and the Moon is not a matter of executing a succession of exploratory missions, no matter how successful they may be individually or collectively. It is a problem of a different nature: a **dynamic logistics regime** whose behavior is governed by the interaction of physical laws, logistical principles, and known technologies—elements that until now have been studied and applied independently. It is the **integration** of these elements into a continuous operational system that gives rise to **emergent behavior**, with its own properties that cannot be deduced from the isolated analysis of each of its components.

STATE OF THE ART (RELATED WORK)

Over the past two decades, interest in space logistics has grown significantly as a consequence of the transition from exploration programs based on isolated missions toward architectures intended to sustain a prolonged human presence in cislunar space. This shift has driven the development of studies on orbital propellant depots, reusable vehicles, in-situ resource utilization (ISRU), permanent infrastructure, optimization of logistics campaigns, and space transportation networks.

The most comprehensive review published to date on this discipline is **Koki Ho's *Space Logistics Modeling and Optimization: Review of the State of the Art*** (2024), which compiles and classifies several decades of research in space logistics. The author identifies as the main lines of work: logistics network modeling, exploration campaign optimization, orbital depot location, inventory management, route planning, and infrastructure design using techniques from Operations Research such as **Network Flow**, **Vehicle Routing Problem (VRP)**, **Facility Location Problem (FLP)**, **Mixed Integer Linear Programming (MILP)**, and stochastic inventory models.

Within this framework, the logistics network constitutes a mathematical representation used to model a previously defined architecture and subsequently optimize its operation. The object of study remains the logistics campaign or the architecture itself—**not the network as an independent physical or mathematical entity**.

A second widely developed line of research corresponds to reusable lunar architectures based on orbital propellant depots. Representative works include ***Reusable Lunar Transportation Architecture Utilizing Orbital Propellant Depots*** (AIAA, 2009), as well as **NASA's *Architecture Study for a Fuel Depot Supplied from Lunar Resources*** (2016). Both analyze multiple configurations of vehicles, depots, and transfer strategies with the goal of minimizing initial mass, cost, number of launches, or operational complexity.

In all these studies, the orbital depot is an element of an assumed architecture whose location and sizing are part of the optimization process. It is not presented as a

necessary consequence of the physical constraints of space transportation, nor are general properties associated with its existence within a logistics regime examined.

A third line of research focuses on the design of reusable vehicles, particularly lunar landers and transport systems between orbit and surface. Works such as ***Analysis and Optimization of Lunar Exploration Architecture Based on Reusable Human Spacecraft***, ***Multidisciplinary Design and Architecture Optimization of a Reusable Lunar Lander***, and various NASA studies on reusable lander modeling employ multidisciplinary methodologies to simultaneously optimize vehicle, mission, trajectories, and operational costs.

Again, the object of study remains the vehicle and its integration within a specific architecture. Reusability appears as a technological property of the transport system, not as a mathematical consequence derived from a general logistics regime.

More recently, NASA itself has begun to develop a vision increasingly oriented toward permanent infrastructure. Initiatives such as **NextSTEP Lunar Logistics and Mobility Studies** or **Space Superhighway** explicitly consider logistics capabilities, surface mobility, orbital refueling, maintenance, assembly, and in-space manufacturing as fundamental elements for sustaining human presence beyond Earth orbit.

These programs represent an important shift from exploration based exclusively on missions. However, they continue to approach logistics as a problem of systems engineering, technological development, and operational capabilities — **without formulating a general theory of the logistics regime governing such infrastructures**.

Of particular interest is **NASA's *A Cis-Lunar Propellant Infrastructure for Flexible Path Exploration and Space Commerce*** (2012), where the concept of **bootstrap**, progressive infrastructure growth, and the emergence of a cislunar economy is explicitly introduced. This work is likely the conceptual approach closest to the framework developed in the present theory.

Nevertheless, even in this case, concepts such as bootstrap, reusability, or infrastructure evolution are introduced as design principles or engineering strategies intended to guide architectural development. **They are not mathematically deduced** from general principles nor derived as necessary properties of the system.

The joint analysis of the reviewed literature reveals a common pattern. Existing works study vehicles, architectures, depots, logistics campaigns, infrastructures, or optimization models. Even when network models are used, they constitute mathematical tools intended to represent and optimize a previously defined architecture.

In none of the reviewed works has a formulation been identified whose objective is to identify, characterize, and mathematically define the Earth–Moon logistics regime as an independent object of study.

As far as this review reaches, no formulation has been found that:

• *defines a general logistics efficiency of a segment independent of any specific architecture;*

• *defines a global efficiency of the logistics regime as a mathematical property of the system;*

• *establishes a set of axioms describing the general behavior of the regime; • proves general theorems about said regime;*

• *derives the topology of the network as a consequence of the physical constraints imposed by the Tsiolkovsky equation and orbital mechanics;*

• *mathematically characterizes the evolution of the logistics regime through general properties independent of the particular vehicles employed.*

Consequently, the present theory does not aim to propose a new lunar architecture nor to optimize an existing exploration campaign. **Its object of study is different. It considers the Earth–Moon Logistics Regime as a mathematical structure whose existence emerges from the interaction between the Tsiolkovsky equation, orbital mechanics, and the general principles of logistics.**

The purpose of the present theory is to identify that object, characterize it, and define it mathematically through definitions, axioms, theorems, and formal models, independently of the particular architectures that may implement it. It does not seek to establish engineering recommendations nor propose a specific technology, but rather to develop a general mathematical framework that allows describing the behavior of the Earth–Moon logistics regime as a physical system subject to its own laws, properties, and invariants.

From this perspective, vehicles, depots, infrastructures, and campaigns cease to constitute the primary object of analysis and instead become particular realizations of a single underlying logistics regime, whose existence and properties constitute the true object of study of the present theory.

FOUNDATIONAL DISCIPLINES

The theory does not introduce new natural laws. It relies on established principles from different disciplines, each of which contributes a specific dimension to the system. It is their integration, and not each element separately, that determines the behavior of the whole.

Physics of transportation (Tsiolkovsky)

The rocket equation imposes the segmentation of Δv as a direct consequence of physics. The Tsiolkovsky equation does not prescribe a logistics regime, but it establishes the boundary conditions that make its existence necessary.

Orbital mechanics

It establishes the Δv required to move within the Earth–Moon environment and, by extension, throughout the Solar System.

Network theory

The segmentation of Δv conditions the location and function of the nodes. The topology of the network emerges from the physical constraints of the system and determines its operational properties.

Logistics (naval, terrestrial, and aerial)

Cargo flows, propellant flows, and vehicle flows respond to different functions and must be managed in a decoupled manner. This principle is common to modern logistics networks and can be transferred to the Earth–Moon system.

Industrial engineering

The functional specialization of vehicles is a consequence of the segmentation of functions and the existence of differentiated flows. The efficiency of the system depends on the specialization of its elements, not on the versatility of each individual vehicle.

Philosophy of technology

Operating under constraints does not consist of maximizing the performance of each component, but of selecting solutions that are sufficiently simple, reliable, predictable, and resilient to fulfill a specific function. The technological frontier does not constitute, by itself, a design criterion.

Dynamic systems

A bidirectional logistics network constitutes a dynamic system with emergent properties. Its operation implies continuous variations in inventories, flows, and resource utilization. The goal is not to reach a static state, but to maintain continuous operation within defined operational limits.

Establishment of standards

The existence of a continuous logistics infrastructure creates the need for open standards for interfaces, operations, and services. Standardization is a necessary condition for scalability and interoperability of the system.

Market economy

Open standards reduce barriers to entry and encourage the incorporation of multiple operators. The development of a competitive market does not precede the logistics infrastructure; it emerges as a consequence of its consolidation.

FROM FOUNDATIONS TO AXIOMS: EPISTEMOLOGICAL STATUS OF THE THEORY

The foundations stated in the previous section are not original propositions nor do they constitute a new contribution to the knowledge of the disciplines from which they come. They do not introduce new laws, do not correct existing laws, and do not propose heterodox reinterpretations of any established body of knowledge.

The contribution of this theory does not consist in proving what is already proven. It consists in integrating these principles—tested separately, but never before articulated as a coherent system—into a single framework that relates them and extracts from that relationship consequences that are not visible when each discipline is considered in isolation.

These principles are an integral part of human knowledge with universal character. We accept as valid the content of all these disciplines: physics of transportation, network theory, naval, terrestrial, and aerial logistics, industrial engineering, dynamic systems theory, the science of standards, and market economics. It is not necessary to prove them. They are already proven.

Therefore, we can extract from them the axioms of this theory. We assign these axioms the truth value “true” by construction, because they are contained and proven in the disciplinary corpora from which they originate. They require no additional proof. They will not be demonstrated within this theory.

This procedure is not arbitrary. It is analogous to the one used by Euclid when establishing the axioms of geometry: statements whose truth is not demonstrated within the system, but assumed as a starting point because they rely on already established knowledge or on the evidence of facts.

The theory, from here on, does not prove its axioms. It declares them. And upon them it builds a system of principles, rules, and operational consequences.

Axiom 1

Derived from: Physics of transportation (Tsiolkovsky)

The segmentation of Δv is a physical constraint, not a design choice.

The relationship between mass and velocity in an impulsive propulsion system is logarithmic. As a consequence, the energetic cost of covering a high Δv in a single stage grows exponentially, making any transportation architecture that does not segment the trajectory into successive stages unviable. Segmentation is not a design option: it is an imposition of physics.

Axiom 2

Derived from: Physics of transportation (Tsiolkovsky) + Network theory

The segmentation of Δv imposes the existence of nodes. Their exact location admits variants, but is not arbitrary.

The division of the trajectory into stages defines transfer points between segments. These points are the nodes of the network. Physics does not determine a single possible node, but it does restrict the space of options: different locations produce topologically equivalent networks with different performances. The choice of node does not change the nature of the system, but it alters the distribution of its operational loads and its overall performance.

Axiom 3

Derived from: Logistics (naval, terrestrial, and aerial)

Cargo and propellant flows are functionally distinct and must be managed in a decoupled manner.

Cargo and propellant have different origins, destinations, frequencies, and volumes. Treating them as a unified flow — as traditional mission architecture does — is a structural inefficiency, not a technological limitation that can be overcome. Decoupling is not a design option: it is a condition for continuous operation.

Axiom 4

Derived from: Industrial engineering

The functional specialization of vehicles is a consequence of the segmentation of the system, not a design preference.

Expecting a single vehicle to cover multiple segments is not a versatile solution: it is a suboptimal solution in terms of performance, in all the segments it traverses. Specialization is one of the pillars that enabled the exponential increase in productivity during the Industrial Revolution; it is not a choice, it is the direct consequence of the previous axioms.

Axiom 5

Derived from: Philosophy of technology

A technical system is evaluated exclusively by the degree to which it fulfills its declared objectives, measured at the level of the system as a whole, not its individual components.

Technological sophistication, maximum component performance, or symbolic meaning do not constitute valid evaluation criteria. A component may be suboptimal in its specific technical dimension and yet contribute to the success of the system as a whole, and that contribution is the only thing that matters. Conversely, a technically superior component that compromises the operability, simplicity, or resilience of the whole is a dysfunction, not an achievement. Each component can and must be evaluated in relation to its specific function within the system. Failure and success are ontologically opposed. There is no such thing as successful failure.

Axiom 5 has a different status from the previous ones. It does not derive from a proven scientific corpus, but from a position on how a technical system is evaluated. Its axiomatic value does not come from physics or mathematics, but from the deliberate decision to adopt this criterion as a foundation of the theory. It is a choice of framework, not a consequence of nature.

Axiom 6 — THE HINGE

Derived from: Dynamic systems

A continuously operating logistics network is a dynamic system with emergent properties that cannot be deduced from the isolated analysis of its components.

When the five previous elements—Δv segmentation, node topology, flow decoupling, vehicle specialization, and operability criterion—are integrated into a system in continuous operation, the whole acquires emergent properties that cannot be deduced from the analysis of each component separately. This axiom marks the boundary between two domains of the theory. The previous axioms describe what physics and technique impose on the system. The following axioms describe what the system produces when it operates.

Axiom 7

Derived from: Standardization and interoperability

The joint operation of vehicles and nodes inevitably imposes a minimum degree of standardization in their interfaces. This standardization is not a design decision — it is an operational consequence of the existence of the network.

As soon as two vehicles share a node, something has to match between them. It is a condition that the network imposes by the mere fact of existing. What does require a conscious decision is the depth of that standardization and its governance. Physics

imposes interface compatibility, but the accessible or restrictive character of that interface responds to different logics that exceed the technical nature of the network. Those implications appear only when the consolidated system interacts with the economic environment.

Axiom 8

Derived from: Market economy

The market does not precede the logistics infrastructure. It emerges as a consequence of its consolidation and of the standardization that makes it operable.

The sequence is invariant: first the infrastructure, then the standards, then the market. Reversing that order is the conceptual error that has paralyzed the development of the cislunar economy. Subsequent development confirms this with historical evidence.

EXPLANATION OF AXIOM 1

The Tsiolkovsky equation, formulated more than a century ago, is the mathematical foundation of all modern astronautics and the reason multistage rockets exist. It is not an elegant engineering solution nor a historical preference: it is the only possible answer to what physics imposes.

The equation relates the change in velocity of a vehicle — Δv — with the initial mass, final mass, and specific impulse of the propellant. The relationship is not linear: it is logarithmic. And that mathematical characteristic has consequences that are not immediately intuitive.

At low Δv, the curve is relatively forgiving: gaining additional velocity requires more mass, but the penalty is manageable. At high Δv, the curve becomes abrupt: each additional meter per second requires a disproportionately larger amount of propellant. The vehicle progressively becomes a fuel container with a residual payload. The consequence is inevitable: no transportation architecture that attempts to traverse the entire Earth–Moon trajectory in a single stage can be efficient. **Segmentation is not a design option**. It is what the equation demands.

But segmenting the trajectory is not a neutral act. Each cut in Δv defines a transfer point between stages. It defines a **segment** of the network. And any concatenation of segments ends in a **node**. A node is not just a stop: it is the place where flows separate, where vehicles specialize, where the system begins to acquire architecture. Tsiolkovsky did not design a logistics network. But he wrote it into his equation. One hundred years later, the consequence remains the same.

PRINCIPLE 1

Derived from: Axiom 1

Δv segmentation not only divides the trajectory. It generates the architecture of the system.

EXPLANATION OF AXIOM 2

Every network has nodes. It is not a choice; it is a structural condition of any network. In the Earth–Moon system, nodes are not invented nor placed for political or historical convenience. They emerge from Δv segmentation. Each cut in the trajectory defines a transfer point between stages, and that point is a node. Physics imposes it before anyone makes any design decision.

But physics does not dictate a single possible node. Orbital mechanics defines a family of options: orbits of different altitudes, different inclinations, different stability. LLO and NRHO are both valid nodes for the cislunar segment. Neither violates physics. But they are not equivalent: the choice between them is not arbitrary — it is restricted by physics — but it is also not unique. It is a technical decision within a space of options that physics defines and logistics orders by performance.

What is non-negotiable is the existence of the node. What admits variants is its exact location. (*See Mathematical Annex II — Numerical Verification, section 2*.)

PRINCIPLE 2

Derived from: Axiom 1 + orbital mechanics

In the Earth–Moon system, the location of nodes is not arbitrary. Orbital mechanics and Δv segmentation restrict the space of options. Within that space, logistics allows ordering the options by their consequences for network performance.

EXPLANATION OF AXIOM 3

Nodes are not just points of geometric transfer. They are the place where the system makes decisions. And the first decision a node imposes is this: not everything that arrives is the same.

A node receives payload, instruments, materials, people. And it receives propellant, the fuel that moves vehicles between nodes. These two flows do not have the same origin, nor the same destination, nor the same frequency, nor the same volume. They do not have the same concept: payload is what is transported, the reason the logistics network exists. Fuel is only a **consumable**, an input.

No human logistics network mixes them. None. Without exception. In logistics, taking an input for a ride does not exist.

A freighter leaving Shanghai for Rotterdam does not carry in its hold the fuel needed for the return trip. It will refuel where the fuel is, in Rotterdam, or somewhere along the route. Transporting the fuel it will burn in three weeks is not an inefficiency: it is an absurdity. It is so obvious that no shipowner would consider it for even an instant.

A cargo plane crossing the Atlantic does not take off with half its capacity occupied by the kerosene for the return flight. Fuel is managed where the fuel is. Cargo is managed where the cargo is. They are different flows and are treated as different flows.

The space industry is the only known exception to this universal rule. Every mission carries its own return propellant. It is the Apollo architecture, brilliant for exploration, structurally unviable for a continuous-operation regime.

The lander of Model 1 (Ms=13.57 t, Mp=38.63 t, ε=0.26, Isp=340 s) delivers 36.73 tons of payload descending from LLO to the lunar surface, with the return propellant produced locally. The same lander, operating in closed cycle and also carrying the return propellant, delivers 11.04 tons. The difference is not technological. It is logistical. It is not in the lander, which is the same. It is in the network. (*See Mathematical Annex I — Analytical Demonstration*.)

In a continuous-operation regime, propellant is managed where the propellant is. Cargo is managed where the cargo is. Mixing both flows is not an operational simplification; it is a **structural deficiency** with quantifiable consequences.

PRINCIPLE 3

Derived from: Axiom 3

Propellant is not payload. It is a consumable. The propellant flow must be decoupled from the cargo flow or accept the inefficiency.

EXPLANATION OF AXIOM 4

A logistics network is a system and operates as such. Therefore it must be designed as such. The parts of this system serve a common objective. It is a known fact applied in industrial design, from cars to facilities. The criterion for sizing a circulation pump is that it delivers the flow rate and backpressure the installation needs. The same principle applies to the components of the logistics network. A vehicle is a solution to a specific problem. The problem is defined by the segment: the Δv it must traverse, the environment in which it operates, the mass it must move, and the frequency with which it must do so.

Ascent from Earth’s surface to low orbit requires high-thrust engines capable of overcoming gravity and atmosphere. Cislunar transit requires energy efficiency in vacuum. Lunar descent requires precise control at low Δv in the absence of

atmosphere. These three functions are not variants of the same function; they are different functions that impose different and mutually incompatible requirements. A lander that must operate for years in cislunar conditions has redundancy and resilience requirements that an upper stage discarded after 3 hours does not have.

A vehicle that attempts to cover two or more segments is not versatile**; it is suboptimal in all the segments it traverses**. It carries unnecessary structural mass in each segment. It compromises its Isp, its structural fraction, or its payload capacity to satisfy requirements that exist only in another segment. This is particularly applicable to logistics networks. **Not necessarily to a mission architecture**. A lander that assumes several functions may be a practical solution in that exploratory context. Like a Swiss Army knife. Its versatility is appreciated in a summer camp, but no one works with it in a machine shop.

Specialization is not an aesthetic preference. It is the correct engineering response to the direct consequences of Δv segmentation, flow decoupling, and the existence of nodes.

The most direct method to parameterize the vehicles that cover the different segments is to apply the Tsiolkovsky equation according to the characteristics of each segment. In this way, each vehicle ensures maximum performance for its particular segment.

PRINCIPLE 4

Derived from: Axioms 1, 2, 3, and 4

Each vehicle or depot in the network must be sized to fulfill the function assigned by its segment, no more, no less. Versatility is not a virtue in a logistics regime. It is a design deficiency.

EXPLANATION OF AXIOM 5

This is the axiom with a different status: it does not derive from physics, but from an epistemological stance. Its axiomatic value does not come from physics or mathematics, but from the deliberate decision to adopt this criterion as a foundation of the theory. This decision is not a whim. It obeys the need to establish a univocal criterion for evaluating success.

This theory establishes the general principles for the study and design of logistics networks in cislunar space. From these principles an indeterminate number of models can be derived, and each model can be implemented in a solution with specific hardware. Inevitably, some of these implementations will become a physical reality.

Throughout this process of transferring the principles to a real and operational logistics network, it is essential to define under which criteria the results will be evaluated. The establishment of clear and measurable objectives is not a methodological suggestion; it is the condition of possibility for the system to acquire technical and economic meaning.

Industrial design knows well the difference between optimizing a component and optimizing a system. A pump that delivers the maximum possible flow rate is not necessarily the correct pump. The correct pump is the one that delivers the flow rate the installation needs. The criterion is not the maximum performance of the component. It is the fulfillment of the function the system assigns to it. The evaluation of each component must be based on its contribution to the success of the whole.

The mathematical appendix adds the numerical demonstration. At the component level, the autonomous vehicle of Case B appears more efficient: it consumes less propellant per cycle and does not require tanker flights. At the system level, it delivers 71.3% less payload with a logistics return rate 50% lower. The component criterion and the systemic criterion not only differ: they point in opposite directions. (*See Mathematical Annex I — Analytical Demonstration*).

This principle, universal in engineering, is systematically violated in the space industry. And it is violated in a particular way: by confusing technological sophistication with success.

In the history of astronautics, we have two paradigmatic examples. The Apollo program set a clear and concrete objective, stated before its beginning and universally known: *"this nation should commit itself to achieving the goal, before this decade is out, of landing a man on the Moon and returning him safely to the Earth".* And every design decision was oriented toward fulfilling this objective. And it succeeded. It fulfilled it. That is success. And the logical consequence of this success was its cancellation. Once the objectives were achieved, there was no reason to continue the project. Because permanence on the Moon was never Apollo's objective.

The Space Shuttle is another paradigmatic case. A machine of extraordinary complexity and technological ambition. Liquid hydrogen engines at the limit of the art. Unique orbital return capability in its time. An incredible machine. And a failure in terms of its declared objectives: reducing the cost of access to space, increasing flight cadence, making orbital access routine. It fulfilled none. The Shuttle was not a successful failure; it was a failure. **And failure and success are ontologically opposed**.

PRINCIPLE 5A

Derived from: Axiom 5

The objectives of a technical system must be clearly stated before proceeding with its design. Success is the degree of fulfillment of those objectives, measured at the level of the system as a whole. Everything else is noise.

PRINCIPLE 5B

Derived from: Axiom 5

The success of a component of the system must be evaluated according to its contribution to the fulfillment of the objectives of the system as a whole.

Everything else is noise. Any isolated optimization of a component that obstructs this end is counterproductive.

EXPLANATION OF AXIOM 6

The five previous axioms describe constraints: what physics imposes, what logistics requires, what engineering resolves. Each one, separately, is a piece of conventional engineering. Nothing in them predicts what happens when they are integrated and the system begins to operate continuously. But a logistics network is a complex system. And every complex system has emergent properties, properties that cannot be deduced from the study of its components because they arise from the interaction among them.

That is what distinguishes a mission from a regime. A mission is executed, completed, and closed. A regime is not completed; it is sustained. And sustaining itself over time introduces a behavior that static design does not contemplate: the properties of the system as such.

Inventories at nodes rise and fall according to vehicle cadence. Buffers absorb delays and propagate them when they are exhausted. Bottlenecks appear that were not in any design blueprint; they arise from the interaction between components that individually work well. The system develops an operating state toward which it tends, not because someone programs it, but because it is the configuration favored by its own internal constraints.

None of these properties exists in the components. They are not in the lander, nor in the tug, nor in the depot. They exist only in the network operating as a network, over time. That is why this axiom is the hinge: it marks the moment when the theory stops talking about parts and begins talking about the whole.

The study and understanding of the emergent properties of a logistics network is one of the declared objectives of this theory. It constitutes a vast field of research, key for the design, implementation, and operation of these networks, because only by understanding their emergent behavior is it possible to make the correct decisions during their operation.

PRINCIPLE 6

Derived from: Axiom 6

A logistics regime is not designed for a final state. It is designed to sustain itself in continuous operation, absorbing its own internal tensions. Like any complex system, it presents emergent properties that its parts did not present.

The design objective is not static stability, but the system’s ability to maintain its operation within limits that arise from its emergent properties, such as oscillations, delays, and relaxation times.

EXPLANATION OF AXIOM 7

The joint operation of vehicles and nodes inevitably imposes a minimum degree of standardization in their interfaces. This standardization is not a design decision; it is an operational consequence of the existence of the network.

As soon as two vehicles share a node, something has to match between them: the docking port of a lander and that of a surface depot, the propellant a tug delivers and the propellant a lander consumes, the communication protocol between vehicle and node. It is not an industrial policy choice. It is geometry, it is chemistry, it is protocol, and it occurs by the mere fact that two elements of the network must interact. No conscious decision produces it; no conscious decision could avoid it without preventing the network from existing.

What does require a conscious decision is everything that exceeds that structural minimum: the depth of the standardization and its governance.

Physics sets a floor — the indispensable compatibility for transfer to occur — but it does not set a ceiling. A standard may stop exactly at that floor, covering only what interoperability requires, or it may deliberately extend in two different directions: toward protocols and specifications (data formats, telemetry, certification criteria for new operators) and toward shared physical components (batteries, engines, filters, and other frequently replaced elements, standardized across different manufacturers)*. Neither extension is indispensable for two specific vehicles to interact, but both determine how many different actors can join the network without friction, and how much its maintenance can be reduced in cost and accelerated once they are integrated.

That extension is not imposed by physics. It is decided by whoever defines the standard, and that decision is not neutral: a standard that stops at the technical minimum leaves room for each operator to resolve the rest on their own, fragmenting the ecosystem into mutually incompatible solutions. A standard that deliberately extends beyond the minimum (documented, published, accessible) is what effectively enables a third party to join the network without having to negotiate case by case with whoever built it first.

Physics imposes interface compatibility. The accessible or restrictive character of that interface responds to a different logic, which exceeds the technical nature of the network and manifests fully when the consolidated system interacts with its economic environment.

*The most cited historical case of this failure is Apollo 13: the lithium hydroxide cartridges for CO_2 removal in the Command Module and the Lunar Module had incompatible geometries (square in one, cylindrical in the other), and the crew had to improvise an adapter with onboard materials to use the Command Module filters inside the Lunar Module. Two systems of the same program, without any shared physical standard beyond the minimum indispensable for the original mission.

PRINCIPLE 7

Derived from: Axiom 7

Any node or vehicle that interacts with other elements of the network is subject to a minimum of standardization in its interfaces. That minimum is not negotiable. What is decided, and decided consciously, is its degree and its openness.

EXPLANATION OF AXIOM 8

The market does not precede the logistics infrastructure. It emerges as a consequence of its consolidation and of the standardization that makes it operable.

No cislunar transportation market can exist before an infrastructure that operates continuously exists. There is no possible competition between operators if there is nothing to operate, nor possible interoperability if there are no nodes or standards with which to interoperate. The market needs a preexisting system upon which to build.

The historical sequence of every continuous-operation network confirms this order without exception. Railroads existed before the competitive railway transportation market. Telecommunications standardization existed before the market of operators who compete over it today. In no case did the market precede the infrastructure. In all cases, it appeared afterward, once the infrastructure was consolidated and standards made it accessible to multiple operators. The creator of TCP/IP was not thinking about Tinder.

Reversing this sequence, expecting the market to finance the infrastructure before it exists, is the conceptual error that has paralyzed the development of the cislunar economy for decades. No investor finances a market that still has nothing upon which to operate.

The sequence is also confirmed by the barrier to entry. Any actor who today wished to operate in cislunar space, without preexisting infrastructure, would have to create their own vehicles, their own facilities, their own development program, their own control centers, their own communication networks, and also build their own long-duration logistics infrastructure before even being able to think about a business. That is not a barrier to entry; it is the total absence of any possible market. A company needs predictable throughput, stable costs, and the ability to secure its investments. None of those three things exists without a prior infrastructure upon which to operate.

PRINCIPLE 8

Derived from: Axioms 7 and 8

A logistics infrastructure with open standards is the condition that allows the emergence of a cislunar market. Without that infrastructure, any economic activity in space depends exclusively on whoever can finance the entirety of their

own chain, which reduces the field to a handful of state actors or corporations of exceptional scale.

CONSEQUENCES OF ESTABLISHING A LUNAR NETWORK WITH OPEN STANDARDS

Note on the epistemological status of this section

The consequences developed below do not have the same status as Axioms 1 to 4, derived from physical and technical corpora already proven. They share the status of Axiom 5: they are a reasoned stance, not a demonstrable necessity. They are offered with the same epistemological honesty: as the reading this theory makes of the probable consequences of a continuous logistics regime, subject to debate, and not as a closed prediction nor as a normative prescription. In particular, the description that whoever controls the technical standard exercises de facto power over the system is a descriptive statement about the dynamics of any network infrastructure — an observable regularity throughout the history of technical standardization, from railroads to the internet — not a recommendation that such control should be exercised in a closed or exclusionary manner. This theory explicitly defends the opposite in Principle 7 and in point 3 of this same section: that an open and deep standard is preferable, precisely because it dilutes that concentration of power instead of consolidating it. Describing a power dynamic does not equate to legitimizing it nor prescribing it.

CONSEQUENCES

Addressing the economic, political, and conceptual impact of a transportation infrastructure might seem, at first glance, like an analytical excess that escapes the technical and formal character of a treatise on logistics networks. However, within the framework of Complex Systems Theory, these implications are not external additions: they constitute the final and most extensive of the system’s emergent properties. Elementary physics determines the geometry of the nodes; engineering resolves the interfaces; but the continuous-operation regime of these variables inevitably predetermines the geopolitical model of the cislunar economy. The sequence that links the velocity differential (Δv) with space macroeconomics is an uninterrupted logical line.

THE SOVEREIGNTY OF THE STANDARD VERSUS THE COLLAPSE OF CLASSICAL SPACE LAW

The chained reading of the preceding axioms leads to an inevitable conclusion that contemporary academic literature tends to avoid for political correctness: the current international legal framework is materially obsolete.

The 1967 Outer Space Treaty was conceived for a scenario of static exploration missions and state monopolies, where States competed for geopolitical prestige through

discrete events. That classical law is based on the fiction that non-territorial appropriation guarantees a neutral space.

The introduction of a continuous-operation regime through a logistics network destroys this premise completely:

- **The obsolescence of non-appropriation:** In a continuous dynamic system, real sovereignty is not exercised through military occupation of geographic coordinates, but through the control of mass and energy flows. Whoever possesses the transfer node and defines the technical specifications of the interface (propellant chemistry, docking geometry, communication protocols) exercises de facto sovereignty that no diplomatic treaty can neutralize.
- **The standard as material law:** Current ideological currents that attempt to legislate space based on doctrinal preferences or bureaucratic utopias ignore the primacy of infrastructure. Cislunar property rights will not arise from a UN consensus; they will emerge as an operational necessity of the network itself to provide legal security for the exchange of flows in storage buffers.
- **The transition toward cislunar realism:** Well-intentioned regulatory frameworks that do not align with this material logistical reality will collapse under their own irrelevance, because the real market demands cost predictability and inventory security, variables that only physical infrastructure can guarantee.

THE MYTH OF DIPLOMATIC PRESERVATION:

The Antarctic parallel versus the Arctic

To understand the inevitable fate of the cislunar economy, it is mandatory to replace ideology with physical variables and energy-friction variables. The geographic parallel between the Arctic and Antarctica exposes the falseness of the idea that the preservation of a territory depends on the good intentions of a written document:

- **The Arctic model (Economic Integration):** From a geographic and ecological standpoint, both poles present extreme similarities. However, the Arctic is a scenario of intense economic activity and commercial sovereign competition due to its relative physical accessibility and its proximity to consolidated global transportation networks.
- **The Antarctic model (The Treaty Mirage):** Antarctica remains limited to a regime of static, scientific, subsidized bases under the protection of the Antarctic Treaty. The official narrative attributes this preservation to the success of international law. It is a diagnostic error. The real difference lies in the logistical inaccessibility of Antarctica compared to the Arctic; the energy cost of access breaks the profitability equation.

The Moon and the rest of the Solar System are currently trapped in the “Antarctic model,” protected only by the immensity of the gravitational well that separates them from Earth. As soon as a logistics infrastructure with open standards reduces the friction

of the velocity differential (Δv) and lowers access costs, the Moon will undergo the same transition as the Arctic. It will cease to be an isolated scientific laboratory and become an extension of humanity's productive structure.

THE GEOPOLITICAL DIMENSION OF THE STANDARD:

Depth and Openness

The design of a standard is not a neutral decision; determining its depth and level of openness carries immediate geopolitical consequences for that newly integrated territory:

- **The closed standard as a tool of hegemony:** If a dominant actor (state or corporate) imposes proprietary and closed interfaces, it restricts access to the cislunar ecosystem. This forces the rest of humanity to act as captive clients or to duplicate infrastructures, fragmenting space into incompatible and militarizable technological blocs.
- **The open standard as a multilateral catalyst:** Conversely, an open and deep standard democratizes space. It allows a medium-sized State, a university institution, or a specialized corporation to compete or collaborate in a specific segment without having to finance the entirety of the transportation system. Openness dilutes the oligopoly of superpowers and transforms infrastructure into a global public good.

DRASTIC REDUCTION OF ENTRY BARRIERS AND FINANCIAL VIABILITY

The widespread perception that the space economy belongs to the realm of utopia derives from an irrational entry toll: the obligation to assume the vertical deployment of launchers, nodes, communications, and regulatory frameworks before operating commercially.

By providing predictable throughput, stable costs, and operational predictability, the network with open standards absorbs the risk of the base infrastructure. This transforms state-scale projects into commercially viable initiatives financed by the market. The current flourishing of emerging initiatives (startups) demonstrates that latent interest exists; the network is the coupling mechanism that transforms that financial expectation into real economic operations, drastically reducing initial capital barriers.

REDEFINITION OF "PERMANENT ESTABLISHMENT"

The history of traditional astronautics associates the concept of permanent establishment with milestones of static infrastructure: habitat plants, manned scientific bases, or government missions sustained by continuous subsidies. This is an incomplete vision and vulnerable to the political fluctuations of state budgets.

The only real "permanent establishment" on the Moon is to incorporate it definitively into humanity's economic activities.

Permanence is not measured in tons of military or scientific hardware deployed on the surface, but in the anchoring of self-sustaining logistical and commercial flows. A scientific base can be evacuated or defunded; a market interconnected by a bidirectional and open logistics network becomes irreversible. The open-standards network does not seek merely to sustain humans in the vacuum; it seeks to fuse cislunar space with Earth's global productive structure.

NOTICE TO THE READER — WHAT THIS DOCUMENT IS AND WHAT IT IS NOT

THIS IS: an axiomatic mathematical theory describing the structure of the Earth–Moon logistical regime — definitions, axioms, theorems, and proofs.

THIS IS NOT: an architecture proposal, a mission study, or a program assessment.

Accordingly, this document does **NOT** include — because it is not its object of study — cost estimates, development timelines, trajectory simulations, or program-risk analyses.

If what you are looking for is an evaluation of a specific architecture, this is not the document that answers that question.

If what you are looking for is the mathematical structure that any such architecture must satisfy, keep reading.

THEORY OF THE EARTH–MOON LOGISTICS REGIME II

To advance from qualitative description toward a more rigorous analytical framework, the following definitions and theorems are introduced. These instruments allow quantifying the propagation of payload mass through a segmented network and precisely comparing the performance of different architectures.

INTRODUCTION: STANDARD DEFINITIONS AND MATHEMATICAL FRAMEWORK

The following are the classical definitions used in aerospace engineering to describe the performance of a propelled vehicle. These expressions are included as a reminder for the reader and constitute the basis upon which the specific analysis of reusable cislunar systems will be developed.

Structural fraction (ε)

The structural fraction is defined as the ratio between the dry mass of the vehicle and the sum of the dry mass and the propellant mass:

$$\varepsilon = M_{dry} / (M_{dry} + M_{prop})$$

where:

• ε = structural fraction (dimensionless, between 0 and 1)

• M_{dry} = dry mass of the vehicle (tanks, engines, structure, systems)

• M_{prop} = propellant mass

Low values of **ε** indicate a lightweight structure compared to the propellant, which implies better performance. This definition is standard in the analysis of rockets and propulsive stages.

Payload fraction (λ)

The payload mass fraction is defined as the ratio between the payload mass and the total mass of the vehicle at the beginning of a maneuver:

$$\lambda = M_payload / M_0$$

where:

• λ = payload fraction (dimensionless, between 0 and 1)

• M_payload = payload mass

• M_0 = total mass at the beginning of the segment (structure + propellant + payload)

This definition is widely used to compare vehicle architectures and propulsion systems. It is important to note that λ is a property of the vehicle, because it relates the payload to the total mass of the vehicle at the beginning of the maneuver.

Tsiolkovsky equation

The relationship between the required velocity increment and the mass ratio of the vehicle is given by the Tsiolkovsky equation:

$$\Delta V = g_0 \times Isp \times \ln(R)$$

where:

• ΔV = required velocity increment

• g_0 = standard gravitational acceleration

• Isp = specific impulse of the engine

• R = mass ratio = m_0 / m_f

This equation establishes that the mass ratio grows exponentially with ΔV. Consequently, the propellant mass required to complete a maneuver increases more and more aggressively as ΔV increases.

PHYSICAL INTERPRETATION

The logarithmic form of the Tsiolkovsky equation has direct implications for the behavior of the system:

- The mass ratio **R** grows exponentially with **ΔV**.
- The propellant mass required per ton of payload increases sharply as **ΔV** becomes larger.
- The payload fraction **λ** decreases with **ΔV**, initially gently and later more sharply.
- Differences in specific impulse **Isp** have limited impact at moderate **ΔV**, but their effect is amplified at high **ΔV**. However, even in that regime, the exponential penalty affects all propellants and the relative advantages remain bounded when the structural fraction **ε** is the same.

These observations help explain why certain intuitions developed in very high-ΔV missions (for example, interplanetary missions) do not translate directly to cislunar environments, where ΔV values are moderate and structural and operational penalties play a determining role.

TRANSITION TO DEFINITION 1 — THE BLACK BOX

A segment is a black box with two boundaries. Nothing more, nothing less.

Through one boundary, a mass enters: **M_start_i** — everything that enters the box. Through the other boundary, a mass exits: **$M_payload_i$** — everything that leaves the box, counted as the logistical flow of the segment.

“Everything” means exactly that: **everything**. Vehicles, consumables, crew, instruments, spare parts, the guitar someone decides to take to the Moon — any physical object that crosses the input or output boundary defines a mass flow between entry and exit.

This distinction — between what physically crosses the boundary and what is counted as segment flow — is clarified later: for a reusable vehicle, its structure physically crosses the boundary but is **not** counted as flow, because the vehicle itself is permanent infrastructure of the network, not cargo in transit.

The concatenation of two or more segments — the coupling between the exit of one segment and the entry of another — represents a **network segment**. The network is the concatenation of all segments that compose it.

This entire theory speaks about that object. In no case do the properties, theorems, or corollaries apply to a vehicle. We are not talking about rockets. We are talking about the **logistical regime of a logistical network**.

The standard definitions of **ε** and **λ** are useful for describing expendable stages and single-mission vehicles — but they describe the vehicle, not the box. The black box we have just defined requires a different operational magnitude, one that explicitly incorporates:

- the structural mass fraction **ε**,
- the Tsiolkovsky mass ratio **R**,
- the accumulated penalty from multiple operational cycles,
- and the need to compare architectures under realistic logistical conditions.

For this reason, **Definition 1** is introduced below. It formalizes the operational magnitude that measures the flow through that box — not the properties of the vehicle that traverses it.

DEFINITIONS

DEFINITION 1 — Efficiency of a segment

The efficiency η_i of a segment ***i*** of a logistics network is defined as the fraction of mass that, departing from the origin node of the segment, reaches the destination node in the form of useful mass (payload):

$$\eta_i = M_payload_i / M_start_i$$

where:

• **M_start_i** is the total mass at the beginning of the segment (vehicle + propellant + payload)

• **$M_payload_i$** is the useful mass delivered at the end of the segment

Observation: This definition establishes the general concept of efficiency of a segment i of a logistics network. It is important to note that η_i is a property of the segment of the logistics network, because it relates the mass that enters the segment with the mass that exits on the other side. It is applicable to any segment of the network, regardless of whether the vehicle is expendable or reusable.

The magnitude η_i is new in the literature: while **λ** describes vehicles, **η_i** describes segments of a logistics network and allows analyzing reusable architectures and complex operational modes. This distinction remains hidden in the cases that follow immediately — in fact, in Theorem 1 it coincides numerically with λ — and becomes visible only starting from Theorem 4, where logistical recursion separates the two magnitudes.

DEFINITION 2 — Total efficiency of a segmented logistics network

For a network composed of ***n*** segments in series, the total mass-conversion efficiency from origin to destination is the product of the efficiencies of each segment:

$$\eta_{total} = \eta_1 \times \eta_2 \times \eta_3 \times \ldots \times \eta_n$$

or, in abbreviated notation:

$$\eta_{total} = \prod \eta_i \text{ (product from i = 1 to i = n)}$$

Observation:

This expression is not a hypothesis nor an approximation. It is an exact identity derived from mass conservation in series systems, where the mass exiting one segment is the same mass entering the next. It requires no additional proof: it is the operational definition of how mass propagates through a segmented logistics network.

DETERMINATION OF SEGMENT EFFICIENCY

Definitions 1 and 2 establish the general concept of segment efficiency and the way individual efficiencies propagate through a logistics network. However, the determination of the efficiency of a segment depends on the mode in which that segment is served.

There are two fundamental ways to provide logistics service between two nodes of a network.

A) Segment served by an expendable vehicle

In an expendable vehicle, all the mass present at the beginning of the segment is part of the logistics flow. The vehicle structure is lost at the end of the operation and therefore must be considered part of the transported mass.

Consequently,

$$\mathbf{M_start_i = Ms_i + Mp_i + M_payload_i}$$

where M_s is the dry mass of the vehicle and M_p is the propellant mass of the vehicle, and the efficiency of the segment can be calculated directly using Definition 1.

The corresponding analytical expression is developed in Theorem 1.

B) Segment served by a reusable vehicle

The vehicle (lander or tug), being reusable, ceases to be part of the logistics flow and becomes infrastructure that, within the framework of this theory, is considered permanent within the system. Consequently, the efficiency of the segment must not be evaluated with respect to the total mass of the vehicle, but with respect to the mass that must be replenished so that the vehicle can begin a new operational cycle.

There are two fundamental modes of operation:

B.1) Closed cycle (CC)

The vehicle transports from the origin node all the propellant necessary to complete the cycle and return to its initial position. Part of the vehicle's capacity is therefore reserved for return propellant, reducing the transportable payload.

This case is developed in Theorem 2.

B.2) Open cycle

The vehicle transports only the propellant necessary to complete the outbound segment. The propellant required for the return is obtained later through another independent logistics operation. Within this regime, two fundamental cases are distinguished:

- **Open cycle without logistics recursion (CA)**
  The resource required for the return is available at the destination node and its acquisition does not require additional operations from the system itself. The most representative case corresponds to refueling through local propellant production (ISRU).
  This regime is developed in Theorem 3.

- **Open cycle with logistics recursion (CS)**
  The resource required for the return must be transported to the destination node through additional trips of the system itself. As a consequence, a logistics recursion appears that modifies the efficiency of the segment.
  This regime is developed in Theorem 4.

THEOREMS — Adopted terminology

Before developing the theorems of the present theory, it is useful to clarify the terminology adopted to describe the different masses that circulate through a logistics network. The distinction between these concepts is essential for correctly interpreting the mathematical developments presented in the following sections.

MAXIMUM TRANSPORT CAPACITY (P)

In the present theory, a distinction is made between the **physical mass** of an object and the **logistical transport capacity** of a vehicle.

Mass magnitudes are represented by the symbol **M** (for example: $\mathbf{M_s}$, $\mathbf{M_p}$, **M_payload**, **M_start**, **M_final**) and describe concrete physical masses that intervene in the Tsiolkovsky equation.

P is defined as the maximum transport capacity of a vehicle during a given segment, expressed as a mass, under the conditions established for the segment and the logistics regime considered.

The magnitudes **P**, $\mathbf{P_{net}}$, $\mathbf{P_{op}}$ and $\mathbf{P_f}$ do not represent different masses, but the use of the **maximum transport capacity** (payload) of the vehicle within the logistics regime. Consequently, they describe the logistical function performed by the transported mass and not its physical nature. The capacity **P** may be allocated indistinctly to the transport of net cargo, operational cargo, or any combination of both.

Net cargo ($\mathbf{P_{net}}$)

($\mathbf{P_{net}}$) is defined as **net cargo**: the mass that constitutes the final objective of the logistics network and that is effectively delivered to the recipient to satisfy a need external to the logistics system itself. Net cargo represents the reason for the network's existence and constitutes the final product of transportation.

Operational cargo ($\mathbf{P_{op}}$)

($\mathbf{P_{op}}$) is defined as operational cargo: the mass that must be transported to guarantee the functioning and sustainment of the logistics network itself.

It includes, among others, propellants, raw materials for ISRU processes (such as carbon and nitrogen), operational consumables and, in more detailed models, spare

parts, lubricants, maintenance components, or any other input necessary to sustain the logistics regime.

In particular, and to facilitate understanding for the reader, the same concept ($\mathbf{P_{op}}$) may alternatively be expressed as ($\mathbf{P_f}$) if this operational cargo corresponds to propellants (fuel).

Relationship between the concepts

The maximum transport capacity of a vehicle may be used entirely or partially to transport net cargo, operational cargo, or both simultaneously.

Consequently:

$$\mathbf{P = \text{Net cargo} + \text{Operational cargo} = (P_{net}) + (P_{op})}$$

THEOREM 1 — Efficiency of a segment served by an expendable vehicle

Statement:

For a segment of a logistics network served by an expendable vehicle, with structural fraction **ε** and mass ratio **R** determined by the Tsiolkovsky equation, the efficiency of the segment is given by:

$$\eta = (1 - R \times \varepsilon) / [\, R \times (1 - \varepsilon) \,]$$

where:

$$\mathbf{R = \exp(\, \Delta v / (Isp \times g_0) \,)}$$

Variables:

Symbol	Description	Units
η	Segment efficiency	Dimensionless
ε	Structural mass fraction of the vehicle*	Dimensionless
R	Tsiolkovsky mass ratio	Dimensionless
Δv	Velocity requirement of the segment	m/s
Isp	Specific impulse of the engine	s
$\mathbf{g_0}$	Standard gravitational acceleration	9.80665 m/s²

*Note: $\varepsilon = M_s / (M_s + M_p)$, that is, the ratio between the dry mass of the vehicle and the sum of the dry mass and the propellant mass.

Proof

We begin with the definition of segment efficiency (Definition 1):

η = M_payload / M_start

where the initial and final mass balances are:

M_start = M_s + M_p + M_payload
M_final = M_s + M_payload

Tsiolkovsky's equation establishes the relationship between these masses through **R**:

M_start / M_final = R
(M_s + M_p + M_payload) / (M_s + M_payload) = R

From the definition of the structural mass fraction ε = M_s / (M_s + M_p), we can express the dry mass in terms of the consumed propellant mass:

M_s = [ε / (1 − ε)] × M_p

Substituting Ms into the Tsiolkovsky relation, grouping terms, and solving for the payload mass M_payload as a function of M_p and the physical constraint R:

M_payload = M_p × (1 − R × ε) / [(R − 1) × (1 − ε)]

To make the algebraic development explicit before the final cancellation, we expand the equation by substituting M_payload and M_s into the original definition of initial mass (M_start = M_s + M_p + M_payload):

M_start = ([ε / (1 − ε)] × M_p) + M_p + (M_p × (1 − Rε) / [(R − 1)(1 − ε)])

Factoring out M_p and simplifying the fractions using a common denominator, the initial mass reduces to:

M_start = M_p × [R / (R − 1)]

Finally, applying Definition 1 (η = M_payload / M_start), the common factor M_p and the denominator (R − 1) cancel cleanly in the quotient, yielding directly the result for the segment:

η = (1 − R × ε) / [R × (1 − ε)]

Q.E.D.

Validity conditions:

• The vehicle is expendable (it does not return to the origin node).

• All propellant is consumed within the segment.

• The vehicle's dry mass is discarded at the end of the segment.

OBSERVATION 1 — Coincidence between λ and η in a segment served by an expendable vehicle

The definitions of the payload fraction λ and the segment efficiency η are, in appearance, equivalent:

$\lambda = M_{payload} / m_0$ $\eta = M_{payload} / M_{start}$

For a segment served by an expendable vehicle, the initial mass of the segment is exactly the total mass of the vehicle at the beginning of the maneuver:

$m_0 = M_{start} = M_s + M_p + M_{payload}$

Therefore, in this particular case:

$\lambda = \eta$

It is important to emphasize that this numerical coincidence does not imply that λ and η are identical concepts. The payload fraction λ is a property of the vehicle: it expresses what portion of the vehicle's initial mass corresponds to payload. In contrast, the segment efficiency η is a property of the segment within the logistical network: it expresses what portion of the mass leaving the origin node arrives at the destination node as payload.

Theorem 1 **does not introduce a quantity different from λ**; rather, it takes the definition of segment efficiency:

$\eta = M_{payload} / M_{start}$

and expresses it in terms of the physical constraints of the segment (Δv, Isp, R) and the structural mass fraction ε of the expendable vehicle:

$\eta = (1 - R \times \varepsilon) / (R \times (1 - \varepsilon))$

When the values of $M_{payload}$ and m_0 are computed in a manner physically consistent with Tsiolkovsky and with the structural mass fraction ε, the numerical value of η matches exactly the value of λ. The equality breaks only when the input data do not satisfy the conditions of Theorem 1 (for example, when the payload or the initial mass come from an architecture that does not consume all propellant or is not strictly expendable).

This situation is analogous to the relationship between kilogram-mass and kilogram-force: in a gravitational field of intensity g both yield the same number, but they represent **distinct concepts**.

OBSERVATION 2 — The Exponential Nature of Δv Cost

Before continuing with the development, it is worth noting why Δv segmentation is not an engineering preference but an unavoidable consequence of the mathematical form of the Tsiolkovsky equation.

The Tsiolkovsky equation can be rewritten as:

$M_{final} / M_{initial} = e^{(-\Delta v / c)}$, with $c = Isp \times g_0$

The term $e^{(-\Delta v / c)}$ represents the fraction of mass that survives after providing a velocity increment Δv. This fraction does not decrease linearly with Δv: it decreases in proportion to what remains at each instant, exactly like radioactive decay or compound interest. Each additional meter per second of Δv does not cost a fixed amount of mass — it costs a fixed percentage of the mass that still remains.

The direct consequence is that two increments of Δv/2 applied in sequence, each starting from a replenished base, are not equivalent to a single increment of Δv:

$e^{(-\Delta v / c)} \neq 2 \times e^{(-(\Delta v/2) / c)} - 1$ instead, $e^{(-\Delta v / c)} = [\, e^{(-(\Delta v/2) / c)} \,]^2$

The mass that survives a large Δv falls multiplicatively, not additively.

Superimposed on this quantity —which decreases exponentially— is the structural fraction ε, which is a constant magnitude: it does not decrease with Δv; it is a fixed design decision of the vehicle. The efficiency of a segment can be written, from Theorem 1, as:

$\eta_{(\Delta v)} = [\, e^{(-\Delta v / c)} - \varepsilon \,] / (1 - \varepsilon)$

This form makes visible a central mechanism of the entire theory: η is the difference between a quantity that decays exponentially with Δv and a constant that does not decay at all. As long as the first remains above the second, the vehicle transports useful payload. The moment the exponential decay crosses below the fixed value of ε, η becomes negative, and the vehicle ceases to exist as a payload transporter — not gradually, but as a threshold (see Corollary 2, Theorem 2, $\varepsilon_{crit} = 1/R$).

This asymmetry —what decays exponentially versus what remains constant— is the ultimate reason why a mission architecture that requires a single vehicle to provide the entirety of the accumulated Δv without replenishment inevitably approaches that threshold as the mission becomes more ambitious. And it is the reason why segmenting Δv through refueling nodes and specialized vehicles —the central object of this theory— is not an incremental optimization, but the only way to keep the system on the correct side of the threshold indefinitely.

THEOREM 2 — Maximum payload of a reusable vehicle in a closed cycle (without refueling)

Statement

Given a reusable vehicle (lander) with structural mass M_s, propellant capacity M_p, specific impulse I_{sp}, operating between two nodes with outbound delta-V Δv_1 (transporting payload) and return delta-V Δv_2 (without payload), and without refueling at the intermediate node, the maximum payload **P** that can be transported per cycle on the outbound leg and return empty to the origin node is given by:

$$P = (M_p - R_1 \times M_{p2}) / (R_1 - 1) - M_s$$

where:

$\mathbf{M_{p2}} = M_s \times (e^{\wedge}(\Delta v_2 / c) - 1)$
$\mathbf{R_1} = e^{\wedge}(\Delta v_1 / c)$
$\mathbf{c} = I_{sp} \times g_0$

Proof

1. Return leg without payload (empty)

To determine the maximum payload of the outbound leg, it is necessary to know beforehand the propellant that must be reserved to complete the return leg.

The vehicle returns from the destination node to the origin node without payload. Its mass at the beginning of the return leg is $M_s + M_{p2}$, where M_{p2} is the propellant required to complete that leg. Its mass at the end of the leg is M_s.

By Tsiolkovsky's equation:

$$\Delta v_2 = c \times \ln((M_s + M_{p2}) / M_s)$$

Solving:

$$e^{\wedge}(\Delta v_2 / c) = (M_s + M_{p2}) / M_s \; M_{p2} = M_s \times (e^{\wedge}(\Delta v_2 / c) - 1)$$

2. Outbound leg with payload

The vehicle departs from the origin node with a total mass:

$$M_0 = M_s + M_p + P$$

where M_p is the total available propellant (full tanks) and P is the payload transported during the outbound leg.

Upon arrival at the destination node, it has consumed M_{p1} during the outbound leg, so its final mass is:

$$M_f = M_s + P + M_{p2}$$

where M_{p2} is the propellant reserved for the return leg.

By Tsiolkovsky's equation applied to the outbound leg:

$$\Delta v_1 = c \times \ln((M_s + M_p + P) / (M_s + P + M_{p2}))$$

3. Mass ratio

Define:

$$R_1 = e^{\wedge}(\Delta v_1 / c)$$

Then:

$$(M_s + M_p + P) / (M_s + P + M_{p2}) = R_1$$

4. Solving for P

Multiply:

$$M_s + M_p + P = R_1 \times (M_s + P + M_{p2})$$

Expand:

$$M_s + M_p + P = R_1 \times M_s + R_1 \times P + R_1 \times M_{p2}$$

Group terms with P:

$$P - R_1 \times P = R_1 \times M_s - M_s + R_1 \times M_{p2} - M_p$$

$$P \times (1 - R_1) = M_s \times (R_1 - 1) + R_1 \times M_{p2} - M_p$$

Multiply by −1:

$$P \times (R_1 - 1) = M_p - R_1 \times M_{p2} - M_s \times (R_1 - 1)$$

Solve:

$$P = (M_p - R_1 \times M_{p2}) / (R_1 - 1) - M_s$$

Conclusion of Theorem 2

Theorem 2 provides a closed and general expression for the maximum payload that a lander can transport in a closed cycle LLO ↔ surface without refueling at the surface node. It depends exclusively on:

- structural mass M_s
- propellant capacity M_p
- outbound and return delta-V (Δv_1, Δv_2)
- specific impulse I_{sp}

Existence condition

The theorem is valid if and only if:

$P \geq 0$

If $P < 0$, the vehicle cannot complete the full cycle with positive payload. In that case, the system requires refueling at the destination node or a reduction in structural mass.

Applicability of the theorem

Although the result is valid for any reusable vehicle operating in a closed cycle between two nodes, in the present work its main application corresponds to the analysis of a lander operating between a lunar orbit (LLO, NRHO, or an equivalent orbit) and the lunar surface.

Observation: Physical interpretation of the result

Theorem 2 reveals a fundamental property of reusable vehicles in closed cycles: before transporting any payload, the vehicle must reserve part of its propellant capacity to guarantee the return to the origin node.

Consequently, the available transport capacity does not depend solely on the outbound leg, but also on the energetic requirements of the return leg.

This dependence constitutes an essential difference with respect to expendable vehicles studied in Theorem 1.

COROLLARY 1 OF THEOREM 2 — Efficiency of a segment served by a reusable vehicle in a closed cycle

Statement For a reusable vehicle operating in a closed cycle, the efficiency of the segment is given by the ratio between the transported payload and the total mass of the logistic flow that must traverse the segment during each cycle. Since the vehicle's structure constitutes an infrastructure that, within the framework of this theory, is considered permanent within the system, such mass is not part of the segment's logistic flow.

Consequently:

$$\eta_{cc} = P / (P + Mp_1 + Mp_2)$$

where:

• **P** is the transported payload.

• Mp_1 is the propellant consumed during the outbound trip.

• Mp_2 is the propellant reserved to complete the return trip.

Proof

By Definition 1, the efficiency of a segment is given by:

$$\eta = M_{payload} / M_{inicio}$$

In a reusable vehicle operating in a closed cycle, the logistic flow of the segment consists solely of the payload and the consumable propellant required to complete the cycle.

The structural mass of the vehicle is not part of this flow, since it constitutes a permanent infrastructure of the system.

By definition of the logistic flow of the segment:

$$M_{inicio} = P + Mp_1 + Mp_2$$

Substituting this expression into Definition 1, we obtain:

$$\eta_{cc} = P / (P + Mp_1 + Mp_2)$$

Q.E.D.

COROLLARY 2 OF THEOREM 2 — Critical structural fraction for a reusable tug in a closed cycle

Statement

Consider a reusable chemical-propulsion space tug, with specific impulse Isp and constant structural fraction ε = Ms/(Ms+Mp), which must complete a closed cycle (for example, LEO → NRHO → LEO) without intermediate refueling, with the same Δv in both legs of the cycle.

Let R = exp(Δv/(Isp x g0)) be the mass ratio corresponding to one leg of the trip.

The maximum payload transportable by this vehicle, according to Theorem 2, is:

P = K × Ms,

where

K = [(1−ε)/ε − R(R−1)] / (R−1) − 1

Existence condition

Since the dry mass must necessarily be positive (Ms > 0) in order to transport a positive payload (P > 0), it is required that K > 0.

If K ≤ 0, the equation would demand Ms = P/K < 0, which has no physical meaning.

Developing the inequality K > 0:

$$\varepsilon_{crit} = 1/R^2$$

This expression defines the maximum critical structural fraction compatible with the physical existence of the vehicle: above this value, no tug with that Δv and Isp can transport positive payload in a closed cycle, regardless of how large it is built.

Case study

For Isp = 450 s and Δv = 4,000 m/s (per leg):
R ≈ 2.47
ε_crit = 1 / 2.47² = 1 / 6.1009 ≈ 0.164

Validity conditions

1. The vehicle operates in a strict closed cycle (Theorem 2): it carries from the origin all the propellant required for the outbound and return legs, with no intermediate refueling.

2. The same Δv is assumed for the outbound and return legs. If the Δv values differ (as may occur due to asymmetries in real orbital transit), the corollary must be generalized by substituting R with independent R_1 and R_2 in the original expression of Theorem 2, before solving for ε.
3. ε, Isp, and Δv remain constant throughout the complete cycle.

Observation

This result does not depend on the scale of the vehicle: K is dimensionless and depends only on ε, Δv, and Isp. Therefore, no increase in size can compensate for a structural fraction above the critical threshold — unlike the open-cycle regime (Corollary 2 of Theorem 3), where efficiency also does not depend on scale but a physical solution exists for any $\varepsilon \in (0,1)$. In a closed cycle, exceeding ε_{crit} is not a performance penalty: it is the physical non-existence of the vehicle as a payload transporter.

THEOREM 3 — Maximum payload of a reusable vehicle in an open cycle (refueling at the destination node)

Statement

Given a reusable vehicle (lander) with structural mass M_s, propellant capacity M_p, and specific impulse I_{sp}, operating between two nodes with an outbound delta-V Δv_1 while transporting payload, and whose propellant for the return leg is obtained through refueling at the destination node, the maximum payload **P** that can be transported per cycle during the outbound leg is given by:

$$P = M_p / (R_1 - 1) - M_s$$

where:

$c = I_{sp} \times g_0$

$R_1 = e^{(\Delta v_1 / c)}$

M_s is the structural mass of the vehicle

M_p is the total propellant capacity of the vehicle

Δv_1 is the delta-V of the outbound leg

Proof

1. Mass at the beginning of descent (LLO)

The vehicle departs from the origin node with a total mass:

$M_0 = M_s + M_p + P$

where M_p is the total available propellant (full tanks) and P is the payload transported during the outbound leg.

2. Mass at the end of descent (surface)

Upon arrival at the destination node, it has consumed all the propellant (M_p) carried for that leg. Its final mass is:

$$M_f = M_s + P$$

3. Tsiolkovsky equation applied to the outbound leg

$$\Delta v_1 = c \times \ln((M_s + M_p + P) / (M_s + P))$$

4. Solving for P

Applying the exponential:

$$(M_s + M_p + P) / (M_s + P) = e^{(\Delta v_1 / c)}$$

Define:

$$R_1 = e^{(\Delta v_1 / c)}$$

Then:

$$(M_s + M_p + P) / (M_s + P) = R_1$$

Multiply by (M_s + P):

$$M_s + M_p + P = R_1 \times (M_s + P)$$

Expand:

$$M_s + M_p + P = R_1 \times M_s + R_1 \times P$$

Group terms with P:

$$P - R_1 \times P = R_1 \times M_s - M_s - M_p$$

$$P \times (1 - R_1) = M_s \times (R_1 - 1) - M_p$$

Multiply by −1:

$$P \times (R_1 - 1) = M_p - M_s \times (R_1 - 1)$$

Solve:

$$P = M_p / (R_1 - 1) - M_s$$

where:

$$R_1 = e^{(\Delta v_1 / c)}$$

$$c = Isp \times g_0$$

Q.E.D.

Conclusion of Theorem 3

Theorem 3 provides a closed and general expression for calculating the maximum payload that a reusable vehicle can transport during the outbound leg when operating in an open cycle.

The expression depends exclusively on:

• structural mass $\mathbf{M_s}$

• propellant capacity $\mathbf{M_p}$

• outbound delta-V $\mathbf{\Delta v_1}$

• specific impulse **Isp**

Existence condition

The theorem is valid if and only if:

P ≥ 0

If **P < 0**, the vehicle cannot transport positive payload on the outbound leg. In that case, the system requires a reduction in the delta-V of the segment, a decrease in structural mass, an increase in propellant capacity, or an improvement in the vehicle's specific impulse.

Validity conditions

• The vehicle is reusable and operates in an open cycle.

• The vehicle begins the outbound leg with full tanks.

• All propellant carried for the outbound leg is consumed during that leg.

• The propellant required for the return leg is obtained through an independent logistical operation at the destination node.

• The vehicle does not carry from the origin node the propellant required for the return leg.

Applicability of the theorem

This theorem applies to any reusable vehicle operating in an open cycle, that is, when the propellant required for the return leg is obtained through an independent logistical operation and does not form part of the logistical flow of the outbound leg.

Within the framework of this theory, its main application corresponds to the segment between a lunar orbit (LLO, NRHO, or another operational orbit) and the lunar surface, when the return propellant is produced locally through ISRU.

COROLLARY 1 TO THEOREM 3 — Efficiency of a segment served by a reusable vehicle in an open cycle

Statement

For a reusable vehicle operating in an open cycle, the efficiency of the segment is given by the ratio between the transported payload and the total logistical flow that must traverse the segment during the outbound leg.

Since the vehicle's structure constitutes infrastructure that, within the framework of this theory, is considered permanent within the system, and the propellant required for the return leg is obtained through an independent logistical operation, neither of them forms part of the logistical flow of the outbound leg.

Consequently:

$$\eta_{ca} = P / (P + M_p)$$

where:

• P is the transported payload

• M_p is the propellant consumed during the outbound leg

Proof

By Definition 1, the efficiency of a segment is given by:

$$\eta = M_{payload} / M_{start}$$

In a reusable vehicle operating in an open cycle, the logistical flow of the outbound leg consists solely of the payload and the propellant consumed during that leg.

The structural mass of the vehicle does not form part of the logistical flow, since it constitutes permanent infrastructure of the system. Likewise, the propellant required for the return leg does not form part of the segment's flow, since it is obtained later through an independent logistical operation.

Therefore:

$$M_{start} = P + M_p$$

Substituting this expression into Definition 1 yields:

$$\eta_{ca} = P / (P + M_p)$$

Q.E.D.

COROLLARY 2 TO THEOREM 3 — Invariance of logistical efficiency η_{ca} under vehicle scaling

Statement

Given a set of reusable vehicles operating in an open cycle that satisfy the following conditions:

• same structural fraction **ε**

• same specific impulse **Isp**

• same outbound delta-V $\mathbf{\Delta v_1}$

then $\mathbf{\eta_{ca}}$ (efficiency of the LLO → surface segment with ISRU at the surface) is **invariant** under changes in vehicle scale; that is, it does not depend on the absolute values of M_s and M_p, but only on their ratio $\mathbf{M_p / M_s}$.

Proof

We begin with Theorem 3:

$$\mathbf{P = M_p / (R_1 - 1) - M_s}$$
$$\mathbf{\eta_{ca} = P / (M_p + P)}$$

Substituting P into the expression for η_{ca}:

$$\mathbf{\eta_{ca} = [M_p / (R_1 - 1) - M_s] / [M_p + M_p / (R_1 - 1) - M_s]}$$

Multiply numerator and denominator by (R_1 − 1):

$$\mathbf{\eta_{ca} = [M_p - M_s \, x \, (R_1 - 1)] / [M_p \, x \, (R_1 - 1) + M_p - M_s \, x \, (R_1 - 1)]}$$

Simplify the denominator:

$$\mathbf{\eta_{ca} = [M_p - M_s \, x \, (R_1 - 1)] / [M_p \, x \, R_1 - M_s \, x \, (R_1 - 1)]}$$

Divide numerator and denominator by M_s:

$$\mathbf{\eta_{ca} = [(M_p / M_s) - (R_1 - 1)] / [(M_p / M_s) \, x \, R_1 - (R_1 - 1)]}$$

Define:

$$\mathbf{k = M_p / M_s}$$

Then:

$$\mathbf{\eta_{ca} = [k - (R_1 - 1)] / [k \, x \, R_1 - (R_1 - 1)]}$$

The resulting expression depends only on **k** and $\mathbf{R_1}$.

Since:

• k is determined by ε (structural fraction)
• R_1 is determined by I_{sp} and Δv_1
then $\mathbf{\eta_{ca}}$ **does not depend on the scale of the lander**.

Q.E.D.

Operational conclusion

In an open-cycle logistical regime, the efficiency of the segment neither improves nor worsens when changing the scale of the vehicle, provided that the structural fraction, specific impulse, and outbound delta-V remain constant.

Consequently, vehicle sizing may respond to operational, reliability, capacity, cadence, or cost criteria without altering the intrinsic logistical efficiency of the segment.

Particular application

Within the framework of this theory, this result applies to transport between a lunar orbit and the surface using a reusable lander when the return propellant is produced locally through ISRU.

This implies that:

- The sizing of the lander does not affect the logistical performance of the system, provided that surface refueling is 100% ISRU.
- The decision to use a larger or smaller lander must be based on other criteria (operational, reliability, launch cadence, etc.), not on segment efficiency.

Numerical verification

Lander	Ms (t)	Mp (t)	P (t)	η_{ca}
Modelo 2	7,16	20,38	19,38	48,73%
Modelo 1	13,57	38,63	36,74	48,73%
Escala ×2	14,32	40,76	38,76	48,74%
Escala ×0,5	3,58	10,19	9,69	48,74%

THEOREM 4 — Efficiency of a segment served by a reusable vehicle in an open cycle, without ISRU, with a surface depot

In this theorem, a *load cycle* is understood as the complete sequence of operations in which a lander descends from the orbital node transporting a payload **P** to the surface, is subsequently refueled, and returns empty to the orbital node, restoring the initial conditions of the system to begin a new cycle.

The refueling trips $\mathbf{V_f}$ are considered ideal fractions of this reference cycle, represented by the mean value $\mathbf{N_{Vf}}$, which expresses the number of refueling trips required per load cycle in steady-state operation. Each refueling trip is physically identical to the load trip, differing only in that the transport capacity **P** is used entirely to carry propellant.

Statement

Given a reusable vehicle (lander) operating between an orbital node (LLO/NRHO) and the lunar surface, with structural mass $\mathbf{M_s}$ and propellant capacity $\mathbf{M_p}$, which in each load trip consumes all its propellant during descent, and which—because local production does not exist (ISRU = 0)—must refuel its ascent propellant through additional trips of the same vehicle that bring propellant from the orbital node to a surface depot, the efficiency of the segment is given by:

$$\eta_{cs} = P / (P + M_p \times (1 + N_{Vf}))$$

where:

$N_{Vf} = M_{p2} / (P - M_{p2})$

and **P** is the **maximum transport capacity** per trip (Theorem 3), and M_{p2} is the propellant required for the vehicle to complete the return leg empty to the orbital node (Theorem 2).

Variables

Symbol	Description	Units
η_{cs}	Segment efficiency (open cycle without ISRU)	Dimensionless
P	Maximum transport capacity (Theorem 3)	t
M_p	Lander propellant capacity	t
M_{p2}	Empty-ascent propellant (Theorem 2)	t
N_Vf	Refueling trips per load cycle	Dimensionless

Proof

1. Maximum transport capacity per trip

Each lander trip—whether transporting payload or propellant—consumes all M_p during descent, as established in Theorem 3:

$P = M_p / (R_1 - 1) - M_s$

This is the maximum mass the lander can deliver to the surface per trip, either as net payload (P_{net}) or as propellant delivered to the surface depot (P_f).

2. Empty-ascent propellant

Any lander that reaches the surface after consuming all descent propellant requires M_{p2} to return empty to the orbital node, according to Theorem 2:

$M_{p2} = M_s \times (e^{(\Delta v_ascent / c)} - 1)$

3. The recursion problem

Since ISRU = 0, this M_{p2} cannot be produced on the surface: it must be extracted from the surface depot, which is supplied by the refueling trips V_f.

But each V_f trip is physically identical to a load cycle descent:

- it consumes all M_p,
- it arrives empty,
- and it also requires M_{p2} to return.

Therefore, the propellant delivered by the $\mathbf{N_{Vf}}$ trips must cover not only the $\mathbf{M_{p2}}$ of the load-cycle lander, but also the $\mathbf{M_{p2}}$ of each of the $\mathbf{N_{Vf}}$ refueling trips.

4. Steady-state propellant balance

Let $\mathbf{N_{Vf}}$ be the number of refueling trips per load cycle.

Total propellant delivered to the depot:

$$\mathbf{N_{Vf} \times P}$$

Total propellant consumed in ascents:

$$\mathbf{(1 + N_{Vf}) \times M_{p2}}$$

The “1” corresponds to the load-cycle lander; $\mathbf{N_{Vf}}$ corresponds to the refueling trips.

In equilibrium:

$$\mathbf{N_{Vf} \times P = M_{p2} \times (1 + N_{Vf})}$$

5. Solving for N_{Vf}

$N_{Vf} \times P - N_{Vf} \times M_{p2} = M_{p2}$

$N_{Vf} \times (P - M_{p2}) = M_{p2}$

$N_{Vf} = M_{p2} / (P - M_{p2})$

6. Total propellant consumed in LLO per cycle

Each trip—load or refueling—consumes $\mathbf{M_p}$ in its descent.

Thus, per load cycle:

$M_{LLO}_propellant = M_p \times (1 + N_{Vf})$

7. Segment definition at the nodes

Although recursion complicates flight-by-flight accounting, the segment has a clean boundary between nodes:

- Everything that enters LLO to sustain a full cycle (P delivered + M_{LLO}_propellant consumed in descents)

versus

- What actually arrives as useful payload at the surface (P).

Applying Definition 1 at that boundary:

$$\mathbf{\eta_{cs} = P / (P + M_{LLO}_propellant)}$$

Substituting the expression from step 6:

$$\mathbf{\eta_{cs} = P / (P + M_p \times (1 + N_{Vf}))}$$

with:

$$\mathbf{N_{Vf} = M_{p2} / (P - M_{p2})}$$

Q.E.D.

Condition of existence

The theorem is valid if and only if $\mathbf{P > M_{p2}}$. If $\mathbf{P \leq M_{p2}}$, a single refueling trip $\mathbf{N_{Vf}}$ cannot deliver enough propellant to cover even the propellant required for its own return to the orbital node. Consequently, the system cannot sustain a stationary regime without local propellant production (ISRU) or without an additional external source of resupply.

Conditions of validity

- The vehicle is reusable and operates between an orbital node and a surface node.
- The vehicle consumes all its propellant during the outbound leg.
- The propellant required for the return leg is obtained exclusively from a depot located at the destination node.
- The surface depot is supplied solely through additional trips of the same type of vehicle.
- The system is analyzed in stationary regime, understood as an equilibrium of the mean propellant balance over successive cycles of cargo and refueling.
- All cargo and refueling trips are considered physically identical, differing only in the nature of the mass transported.

Applicability of the theorem

The present theorem determines the efficiency of a segment served by a reusable lander in open cycle when no local propellant production exists (ISRU = 0) and the propellant required for the return leg must be transported through additional trips of the same system.

Its application is of particular interest in the analysis of Earth–Moon logistical architectures that employ a surface propellant depot supplied from an orbital node (LLO or NRHO).

Generality of Theorem 4

The present theorem has been developed using as an example a reusable lander operating between an orbital node and the lunar surface, as it constitutes the most intuitive and easily understood case.

However, the mathematical development does not depend on the nature of the vehicle nor on the specific segment considered, but solely on the existence of a reusable vehicle that must transport, through additional trips of the same system, the resources required to complete its operational cycle.

Consequently, the theorem is directly generalizable to any segment of a logistical network that satisfies the established validity conditions. The generalization requires only redefining the physical meaning of the variables and nodes considered, without modifying the mathematical development or the expressions obtained.

Conclusion of Theorem 4

In the absence of local propellant production (ISRU = 0), the need to transport the propellant required for the return leg introduces a logistical recursion that increases the total mass that must traverse the segment per operational cycle.

As a consequence, the efficiency of the segment ceases to depend solely on the transport capacity of the vehicle and the propellant consumed during the outbound leg, and begins to depend also on the mean number of refueling trips required to sustain the system in stationary regime.

Therefore, the expression obtained constitutes the general efficiency of a segment served by a reusable vehicle in open cycle without local propellant production.

Final observation on Theorem 4 — The network as a mathematical object

Theorem 4 is not a mere application of the Tsiolkovsky equation to a more complex case. It is the formal demonstration that the relevant efficiency in a reusable logistical regime does not reside in the individual vehicle, but in a higher-order mathematical object: **the network**.

While **λ** describes the payload fraction of an isolated rocket, **η** (Definition 1) describes the fraction of mass that, departing from one node, arrives as useful payload at the next. Theorem 4 shows that, in the presence of logistical recursion, this segment efficiency can only be correctly defined and calculated within the framework of the complete network (Definitions 1 and 2).

This transition —from the vehicle to the network— is the first result in which the conceptual core of the present theory becomes explicit. From this point on, the object of study ceases to be the rocket and becomes the logistical regime as a dynamical system.

COROLLARY 1 TO THEOREM 4 — Recovery of the non-recursive regime

Statement

When logistical recursion disappears ($N_{Vf} = 0$), the expression obtained in Theorem 4 immediately reduces to the general expression for the efficiency of a segment without recursion:

$$\eta = P / (M_p + P)$$

Proof

Substituting $N_{Vf} = 0$ into the expression of Theorem 4:

$$\eta_{cs} = P / (P + M_p \times (1 + N_{Vf}))$$

$$\eta_{cs} = P / (P + M_p)$$

Q.E.D.

PROPERTIES OF THE LOGISTICAL REGIME

DETERMINISTIC CHARACTER OF THE MATHEMATICAL MODEL

Statement

The Earth–Moon logistical regime, defined by the axioms, definitions, and theorems developed in the present theory, constitutes a **discrete and deterministic dynamical system**.

Its evolution is completely determined by the initial conditions of the system and by the physical and mathematical relations that govern the mass flows between the different nodes of the network.

A system is said to be **deterministic** when its temporal evolution is completely specified by the initial conditions and by the evolution rules that govern it. Consequently, given a set of initial conditions and operating parameters, the future state of the system is uniquely determined.

In particular, all mass flows — from injection in LEO to delivery of payload on the lunar surface, propellant supply, and return of useful payload — are governed by:

- The Tsiolkovsky equation, which relates velocity increment to mass ratio.
- The definitions of segment efficiency (Definition 1) and total network efficiency (Definition 2).
- Theorems 1 through 6, which provide closed-form expressions for mass, propellant, and payload variables as functions of the physical parameters of the system.
- The mass-conservation equations at each node, which establish the balance between incoming and outgoing flows.

Consequently, given a set of initial conditions and operating parameters, the future state of the system is uniquely determined within the framework of the continuous mathematical model.

The variables that describe the system at this level are continuous variables: masses, efficiencies, mean trip ratios, etc. All of them are deterministic functions of the input parameters.

COMPLEXITY IS NOT INDETERMINACY

The existence of multiple vehicles, depots, trajectories, and simultaneous flows does not make the system indeterminate. On the contrary, the complexity of the regime arises from the interaction of a set of completely deterministic relations.

The fractional quantities that appear in the analytical developments —for example, the mean number of refueling trips per cycle or the mean number of launches required to sustain a given logistical flow— do not represent uncertainty or numerical approximations. They represent mean values of a discrete regime.

The physical events that constitute the operation of the network (launches, flights, transfers, or refueling operations) are necessarily discrete. However, the global behavior of the regime can be described through continuous mean flows obtained from the mass balance of the system.

Consequently, the appearance of fractional values is a natural property of the logistical regime and not a limitation of the mathematical model.

From continuity of the model to discreteness of the operation

The mathematical model described by the theorems of the present theory operates with continuous variables: masses, propellants, trip ratios, etc. This allows obtaining closed-form expressions and invariant relations that characterize the stationary regime.

However, the real operation of the network consists of discrete events: each trip is a unit event. There is no such thing as 0.95 trips. A trip either occurs or does not occur. This discretization is a fundamental property of the system, not an approximation.

This duality —continuous model, discrete operation— is analogous to what is found in other dynamical systems:

• In statistical mechanics, Newton’s equations are continuous and deterministic, but the behavior of a gas is described through statistical averages. • In queueing theory, mean flows are continuous, but events (arrivals, services) are discrete. • In terrestrial logistics, truck flows are averaged over time, but each truck is a discrete event.

The present theory is situated in this same framework: the model is continuous and deterministic; the operation is discrete and oscillatory around the mean values of the model.

STATIONARY REGIME

Statement

The Earth–Moon logistical regime does not evolve as a succession of independent missions, but as a permanent system of mass transport between nodes.

Once its operation begins, the network operates in a **stationary regime**, characterized by the existence of **mean flows of mass**, **propellant**, and **payload** that **remain stable**

over time, even though the individual operations that produce them are necessarily discrete.

The term *stationary* is used here in the sense of dynamical systems theory: a system is in a stationary regime when its **mean variables** do not depend on time, even though its **instantaneous variables** may fluctuate.

Stationary regime does not mean immobile state

The stationary regime described in the present theory **does not imply** that depots remain at constant mass nor that all individual cycles are identical.

On the contrary:

- launches occur at discrete instants;
- flights follow with different cadences;
- depots continuously receive and deliver mass;
- buffers (payload and propellant depots) experience oscillations around their mean values during normal network operation.

The stationary character of the regime must be understood in a **macroscopic sense**: the mean variables of the system remain bounded and their global balances are conserved over time. Its operation is not restricted to the systematic repetition of a fixed sequence of operations nor to an invariant cycle, but evolves within **a bounded region of the state space**, maintaining constant its global properties and mean flows.

Consequently, the network operates in a dynamic equilibrium, not in a static equilibrium.

Steady-state with respect to what?

The answer is not immediate. A logistics regime **does not remain steady** with respect to the number of launches, the number of flights, or the instantaneous content of depots. All those parameters fluctuate naturally during network operation.

The steady-state nature of the regime must be sought in another magnitude: **the relations between the different logistics flows that sustain the network**.

Consequently, the regime is not steady with respect to the absolute number of trips, but with respect to the relation between the number of payload trips $N_{(Vp)}$ and the number of refueling trips $N_{(Vf)}$. This relation is the magnitude that characterizes the permanent functioning of the network.

The separation of payload and propellant flows, established in Principle 3, allows defining a characteristic relation between them; this relation constitutes the fundamental parameter that describes the steady-state regime of the network.

The characteristic relation of the regime is defined as:

$$\mathbf{N_{(Vp)} / N_{(Vf)}}$$

where:

- $\mathbf{N_{(Vp)}}$ represents the number of trips dedicated to payload transport
- $\mathbf{N_{(Vf)}}$ represents the number of trips dedicated to transporting the operational resources required to sustain the functioning of the network

This relation constitutes the fundamental parameter that describes the steady-state regime of a segment of the network and expresses the equilibrium between the payload flow and the operational resource flow.

Its value does not depend on the absolute number of trips performed, but on the operational needs imposed by the logistics regime considered.

DETERMINATION OF THE RELATION $N_{(Vp)}$ / $N_{(Vf)}$

The objective of this section is to analytically determine the relation between the number of payload trips $N_{(Vp)}$ and the number of refueling trips $N_{(Vf)}$, as it constitutes the fundamental parameter that characterizes the steady-state regime of the network.

The development will be carried out initially on the simplest case compatible with the hypotheses of the theory. A unidirectional logistics network will be considered, in which useful payload is transported from an orbital node (LLO or NRHO) to the lunar surface by means of a reusable lander. A logistics regime with local propellant production sufficient to fully supply the lander's return leg (ISRU = 100%) will be assumed, corresponding to the case developed in Theorem 3.

To facilitate understanding, the reader may visualize this network as an architecture formed by an Upper Stage that performs injection toward TLI and a Tug that performs orbital insertion (LOI), feeding an orbital node from which a reusable lander operates. This representation is exclusively illustrative and does not restrict the general validity of the mathematical development that follows.

Consequently, to determine the relation $N_{(Vp)}/N_{(Vf)}$, it is sufficient to know the propellant requirements of the lander that operates between the orbital node and the lunar surface.

More generally, the relation $N_{(Vp)}/N_{(Vf)}$ is determined by the propellant demand of the segment served by the refueling system.

THEOREM 5 — Characteristic relation of the stationary regime $N_{(Vp)}/N_{(Vf)}$

Definition of the flows

In a network with flow decoupling (Principle 3), the segment between the LEO node and the LLO node —where the orbital Depot acts as a buffer for payload and propellant— presents two fundamental types of trips:

- $\mathbf{V_p}$: trips dedicated to transporting net payload from LEO to LLO.

● $\mathbf{V_f}$: trips dedicated to transporting operational resources (propellant) from LEO to LLO.

The $Depot_{LLO}$ receives both flows. Net payload ($\mathbf{P_{net}}$) is stored in the payload buffer of $Depot_{LLO}$ awaiting transport to the SUP node by the lander's trips. Propellant ($\mathbf{P_f}$) is stored in the propellant buffer of $Depot_{LLO}$ to be consumed by the lander's descents.

The stationary regime of a segmented logistical network is not defined by the absolute number of trips, nor by the instantaneous content of the depots, nor by the annual number of launches. All these parameters fluctuate naturally during network operation.

The stationary character of the regime must be sought in another magnitude: **the relation between the logistical flows that sustain the network**. More specifically, the relation between the number of cargo trips $\mathbf{N_{(Vp)}}$ and the number of refueling trips $\mathbf{N_{(Vf)}}$ constitutes the fundamental parameter that describes the stationary regime of this segment of the network:

$$\mathbf{N_{(Vp)} / N_{(Vf)}}$$

This relation expresses the equilibrium between the payload flow and the operational-resource flow in the LEO → LLO segment. Its value does not depend on the absolute number of trips performed, but on the operational needs imposed by the segment that this upstream segment supplies.

Statement

For a logistical system composed of:

1. A served segment (LLO → SUP) operated by a reusable lander in an open cycle with ISRU = 100% on the surface, whose properties are:
   - **P_lander**: payload delivered to SUP per descent trip.
   - $\mathbf{M_p}$**_lander**: propellant consumed from $Depot_{LLO}$ during descent per trip.
2. A supplying segment (LEO → LLO) operated by a set of vehicles (upper stage, tugs, etc.), which we will generically call ***tug*** for expository purposes, transporting both net payload (V_p) and propellant (V_f) to $Depot_{LLO}$, with capacity P_{tug} per trip.

the characteristic relation of the stationary regime $N_{(Vp)}/N_{(Vf)}$ for the supplying segment is given by:

$$\mathbf{N_{(Vp)} / N_{(Vf)} = P_lander / M_p_lander}$$

Proof

1. Definition of the flows in the supplying segment (LEO → LLO)

The set of vehicles that compose the supplying segment (which we will generically call *tug*) performs two types of trips from LEO to LLO:

- $\mathbf{V_p}$: trips that transport net payload, with capacity Ptug per trip, stored in the payload buffer of $Depot_{LLO}$.

- **V_f**: trips that transport propellant, with capacity Ptug per trip, stored in the propellant buffer of $Depot_{LLO}$.

Observation: The transport capacity **P_{tug}** is the same for both types of trips, since the vehicles of the supplying segment transport net payload or propellant with the same capacity per trip.

Consequently:

$P_{net} = P_{tug}$

$P_f = P_{tug}$

2. Relation between Vp (tug) and V_p_lander (served segment)

Let $N_{(Vp)}$ be the number of cargo trips of the supplying segment.

Each V_p trip of the supplying segment delivers P_{tug} tons of net payload to $Depot_{LLO}$.

Each V_p_lander consumes P_lander tons of net payload from the payload buffer of $Depot_{LLO}$ (the amount delivered to SUP).

Therefore, the total number of V_p_lander that can be sustained by the set of Vp trips is:

$N_{(Vp_lander)} = [N_{(Vp)} \times P_{tug}] / P_lander$

3. Propellant consumed in $Depot_{LLO}$ by the V_p_lander

Each V_p_lander consumes M_p_lander tons of propellant from the propellant buffer of $Depot_{LLO}$.

Therefore, the total propellant consumed in $Depot_{LLO}$ by all V_p_lander sustained by the Vp trips is:

$M_consumed_LLO = N_{(Vp_lander)} \times M_p$_lander

Substituting:

$M_consumed_LLO = [N_{(Vp)} \times P_{tug} / P_lander] \times M_p$_lander

4. Propellant delivered to $Depot_{LLO}$ by the V_f trips

Each V_f trip of the supplying segment delivers P_{tug} tons of propellant to the propellant buffer of $Depot_{LLO}$.

Therefore, the total propellant delivered to $Depot_{LLO}$ by all V_f trips is:

$M_delivered_LLO = N_{(Vf)} \times P_{tug}$

5. Stationary-regime equilibrium condition

In stationary regime, the propellant consumed in $Depot_{LLO}$ must equal the propellant delivered to $Depot_{LLO}$:

M_consumed_LLO = M_delivered_LLO

Substituting:

[$N_{(Vp)}$ x P_{tug} / P_lander] x M_p_lander = $N_{(Vf)}$ x P_{tug}

6. Solving for the relation $N_{(Vp)}/N_{(Vf)}$

Divide both sides by P_{tug} (> 0):

[$N_{(Vp)}$ / P_lander] x M_p_lander = $N_{(Vf)}$

Rearrange:

$N_{(Vp)}$ / $N_{(Vf)}$ = P_lander / M_p_lander

7. Final expression

$$N_{(Vp)} / N_{(Vf)} = P_lander / M_p_lander$$

Q.E.D.

CONDITIONS OF VALIDITY

Theorem 5 is valid under the following conditions:

1. **ISRU = 100% on the lunar surface.** The lander operates in an open cycle with refueling at destination. The propellant required for the lander's ascent from SUP to LLO is produced locally. Consequently, the served segment (LLO → SUP) does not require V_f_lander trips from LLO. $Depot_{LLO}$ only needs to replenish the propellant consumed in descents. This is the regime defined in Theorem 3, where the lander consumes all of M_p_lander during descent and does not transport from LLO the propellant required for ascent.
2. **The supplying segment (LEO → LLO) is operated by a set of vehicles that transport both net payload (V_p) and propellant (V_f) to $Depot_{LLO}$**. The transport capacity P_{tug} is the same for both types of cargo.
3. **The system operates in stationary regime.** The mean propellant balance in $Depot_{LLO}$ is conserved over time. Discrete oscillations around the mean value are absorbed by the $Depot_{LLO}$ buffers.

OBSERVATIONS

1. The bridge between segments

The relation $N_{(Vp)}/N_{(Vf)}$ = P_lander / M_p_lander constitutes the **bridge between two parts of the network**:

• **The served segment (LLO → SUP):** defines the values of P_lander and M_p_lander based on its own physical characteristics (Theorem 3). This segment *imposes* the

relation, by determining the demand for net payload and propellant that must be satisfied from $Depot_{LLO}$.

• **The supplying segment (LEO → LLO):** implements the relation through the proportion in which it distributes its trips between V_p and V_f toward $Depot_{LLO}$. This segment *executes* the relation, by satisfying the demand imposed by the served segment.

In this sense, the served segment imposes the demand; the supplying segment satisfies it. The characteristic relation of the regime is the link connecting both, through $Depot_{LLO}$ acting as an intermediate buffer.

2. Independence from P_tug

The transport capacity of the supplying segment **does not appear in the final expression**. It cancels out during the proof.

This implies that the characteristic relation of the stationary **regime depends exclusively on the properties of the served segment** (P_lander and M_p_lander), not on the size of the vehicles that compose the supplying segment.

3. Physical interpretation

The ratio $N_{(Vp)}/N_{(Vf)}$ expresses the proportion in which the supplying segment distributes its trips between net-payload transport and propellant transport, both directed to $Depot_{LLO}$.

This proportion is determined by the propellant demand of the served segment: the relation between what the lander delivers (P_lander) and what it consumes (M_p_lander) in each descent trip.

4. Relation to the efficiency of the served segment

The efficiency of the segment served by the lander in open cycle (Theorem 3, Corollary 1) is:

$$\eta_{ca} = P_lander / (P_lander + M_p_lander)$$

Whereas the characteristic relation of the regime is:

$$N_{(Vp)} / N_{(Vf)} = P_lander / M_p_lander$$

Both expressions are related but not identical. The first measures the mass efficiency of the segment; the second measures the proportion of flows in the stationary regime of the supplying segment.

DISCRETE CHARACTER OF THE LOGISTICAL REGIME

The relation $N_{(Vp)}/N_{(Vf)}$ is a mean value. In real network operation, trips are discrete and not fractional. Therefore, the system does not execute $N_{(Vp)}/N_{(Vf)}$ exactly each year, but

oscillates around this value, alternating integer combinations of V_p and V_f that, on average, tend toward the characteristic relation.

These oscillations are absorbed by the $Depot_{LLO}$ buffers (payload capacity and propellant capacity), which act as dynamic dampers. The stationary regime is not an immobile state, but a dynamic equilibrium maintained by the interaction between discrete flows, finite buffers, and a fixed mean proportion.

GENERALITY OF THEOREM 5

Although the proof has been developed using a lander as the served segment and a set of vehicles (tugs) as the supplying segment, the mathematical development does not depend on the nature of the vehicles nor on the specific segments considered, but solely on the existence of:

- **A served segment** that consumes operational resources (propellant) from a depot to perform its function.
- **A supplying segment** that transports both net payload and operational resources to that depot.

Consequently, the theorem is directly generalizable to any segmented logistical network that satisfies the established validity conditions. The generalization requires only redefining the physical meaning of the variables and the segments considered, without modifying the mathematical development or the expressions obtained.

THEOREM 6 — Characteristic relation of the stationary regime $N_{(Vp)}/N_{(Vf)}$ for the LEO → LLO segment, without ISRU on the surface

Definition of the flows

In a network with flow decoupling (Principle 3), the segment between the LEO node and the LLO node —where the orbital Depot acts as a buffer for payload and propellant— presents two fundamental types of trips:

- **V_p:** trips dedicated to transporting **net payload** from LEO to LLO.
- **V_f:** trips dedicated to transporting **operational resources (propellant)** from LEO to LLO.

$Depot_{LLO}$ receives both flows. Net payload (P_{net}) is stored in the payload buffer of $Depot_{LLO}$ awaiting transport to the SUP node by the lander's trips. Propellant (P_f) is stored in the propellant buffer of $Depot_{LLO}$ to be consumed by the lander's descents.

The relation between the number of cargo trips $N_{(Vp)}$ and the number of refueling trips $N_{(Vf)}$ constitutes the fundamental parameter that describes the stationary regime of this segment of the network:

$N_{(Vp)}$ / $N_{(Vf)}$

This relation expresses the equilibrium between the payload flow and the operational-resource flow in the LEO → LLO segment. Its value does not depend on the absolute number of trips performed, but on the operational needs imposed by the segment that this upstream segment supplies.

Statement

The stationary regime of a segmented logistical network is not defined by the absolute number of trips, nor by the instantaneous content of the depots, nor by the annual number of launches. All these parameters fluctuate naturally during network operation.

The stationary character of the regime must be sought in another magnitude: **the relation between the logistical flows that sustain the network**.

For a logistical system composed of:

1. **A served segment (LLO → SUP)** operated by a reusable lander in open cycle without ISRU on the surface, where the propellant required for the lander's ascent from SUP to LLO must be transported from $Depot_{LLO}$ through additional trips of the lander itself. **The properties of the lander are:**

   • **P_lander:** payload delivered to SUP per descent trip (V_p_lander).

   • **M_p_lander:** propellant consumed from $Depot_{LLO}$ during descent per V_p_lander.

   • **M_{p2}:** propellant required for the lander to ascend empty from SUP to LLO.

   • **N_V_f_lander:** number of lander refueling trips (from LLO to SUP, transporting propellant) required to sustain one cargo trip.

According to Theorem 4:

$N_V_f_lander = M_{p2} / (P_lander - M_{p2})$

2. **A supplying segment (LEO → LLO)** operated by a set of vehicles (upper stage, tugs, etc.), which we will generically call *tug* for expository purposes, transporting both net payload (V_p) and propellant (V_f) to $Depot_{LLO}$, with capacity P_{tug} per trip.

The characteristic relation of the stationary regime $N_{(Vp)}/N_{(Vf)}$ for the supplying segment is given by:

$$\mathbf{N_{(Vp)} / N_{(Vf)} = P_lander / [M_p_lander \times (1 + N_V_f_lander)]}$$

Proof

1. Definition of the flows in the served segment (LLO → SUP) without ISRU

In the absence of ISRU, the lander operating between LLO and SUP must:

1. **Descend** from LLO to SUP transporting net payload (V_p_lander). For this, it consumes M_p_lander from $Depot_{LLO}$.

2. **Ascend** from SUP to LLO, empty or with payload, consuming M_{p2} tons of propellant. This propellant is not produced locally; it must be transported from LLO to SUP through refueling trips of the lander itself (V_f_lander).

According to Theorem 4, the number of lander refueling trips required per cargo trip is:

$N_V_f_lander = M_{p2} / (P_lander - M_{p2})$

2. Total propellant consumed from $Depot_{LLO}$ by the served segment

For each cargo trip of the lander (V_p_lander), $Depot_{LLO}$ must supply:

- **M_p_lander:** for the descent of the cargo trip.
- **M_p_lander × N_V_f_lander:** for the descents of the lander's refueling trips (each V_f_lander consumes M_p_lander in its own descent).

Therefore, the total propellant consumed from $Depot_{LLO}$ per V_p_lander is:

$M_p_per_V_p_lander = M_p_lander × (1 + N_V_f_lander)$

3. Relation between V_p (tug) and V_p_lander (served segment)

Let $N_{(Vp)}$ be the number of cargo trips of the supplying segment.

Each V_p trip of the supplying segment delivers P_{tug} tons of net payload to $Depot_{LLO}$.

Each V_p_lander consumes P_lander tons of net payload from the payload buffer of $Depot_{LLO}$.

Therefore, the total number of V_p_lander that can be sustained by the set of V_p trips is:

$N(V_p_lander) = [N_{(Vp)} × P_{tug}] / P_lander$

4. Total propellant consumed in $Depot_{LLO}$ by the V_p_lander

Each V_p_lander consumes M_p_lander × (1 + N_V_f_lander) tons of propellant from the propellant buffer of $Depot_{LLO}$ (expression 2).

Therefore, the total propellant consumed in $Depot_{LLO}$ by all V_p_lander sustained by the V_p trips is:

$M_consumed_LLO = N(V_p_lander) × M_p_lander × (1 + N_V_f_lander)$

Substituting:

$M_consumed_LLO = [N_{(Vp)} × P_{tug} / P_lander] × M_p_lander × (1 + N_V_f_lander)$

5. Propellant delivered to $Depot_{LLO}$ by the V_f trips

Each V_f trip of the supplying segment delivers P_{tug} tons of propellant to the propellant buffer of $Depot_{LLO}$.

Therefore, the total propellant delivered to $Depot_{LLO}$ by all V_f trips is:

$M_delivered_LLO = N_{(Vf)} \times P_{tug}$

6. Stationary-regime equilibrium condition

In stationary regime, the propellant consumed in $Depot_{LLO}$ must equal the propellant delivered to $Depot_{LLO}$:

M_consumed_LLO = M_delivered_LLO

Substituting:

$[N(Vp) \times P_{tug} / P_lander] \times M_p_lander \times (1 + N_V_f_lander) = N_{(Vf)} \times P_{tug}$

7. Solving for the relation $N_{(Vp)}/N_{(Vf)}$

Divide both sides by P_{tug} (> 0):

$[N_{(Vp)} / P_lander] \times M_p_lander \times (1 + N_V_f_lander) = N_{(Vf)}$

Rearranging:

$N_{(Vp)} / N_{(Vf)} = P_lander / [M_p_lander \times (1 + N_V_f_lander)]$

8. Final expression

$N_{(Vp)} / N_{(Vf)} = P_lander / [M_p_lander \times (1 + N_V_f_lander)]$

Where:

$N_V_f_lander = M_{p2} / (P_lander - M_{p2})$

Q.E.D.

OBSERVATIONS

1. Physical interpretation

The ratio $N_{(Vp)}/N_{(Vf)}$ expresses the proportion in which the supplying segment distributes its trips between net-payload transport and propellant transport.

In the absence of ISRU, the served segment requires **more propellant from LLO** than in the ISRU case, because it must supply both the descents of the cargo trips and the descents of the lander's refueling trips. This increase is reflected in the factor (1 + N_V_f_lander), which reduces the value of $N_{(Vp)}/N_{(Vf)}$.

2. Relation with Theorem 4

Theorem 4 establishes the efficiency of the served segment (LLO → SUP) without ISRU:

$\eta_{cs} = P_lander / [P_lander + M_p_lander \times (1 + N_V_f_lander)]$

Theorem 6 establishes the characteristic relation of the regime for the supplying segment:

$N_{(Vp)} / N_{(Vf)} = P_lander / [M_p_lander \times (1 + N_V_f_lander)]$

Both expressions share the same denominator. The relation between them is:

$N_{(Vp)} / N_{(Vf)} = \eta_{cs} / [1 - \eta_{cs}]$

3. The role of M_{p2} in Theorem 6

The reader will observe that the final expression of Theorem 6:

$N_{(Vp)} / N_{(Vf)} = P_lander / [M_p_lander \times (1 + N_V_f_lander)]$

does not explicitly contain the term M_{p2} (propellant required for the lander's ascent from the surface to LLO).

This apparent absence is not an omission. The explanation is as follows:

1. The propellant balance on the surface (Theorem 4) determines the number of lander refueling trips required per cargo trip:

$N_V_f_lander = M_{p2} / (P_lander - M_{p2})$

This expression incorporates M_{p2} and P_lander as its only parameters. Therefore, M_{p2} is implicitly included in the value of N_V_f_lander.

2. The balance at DepotLLO accounts only for the propellant that leaves LLO toward SUP, that is, the propellant consumed in descents. The propellant consumed in ascents (SUP → LLO) is not accounted for at LLO, because it is consumed on the surface.
3. What *is* accounted for at LLO is the propellant that must be transported to the surface to enable those ascents. That propellant is exactly:

$P_lander \times N_V_f_lander = M_{p2} \times (1 + N_V_f_lander)$

And it is delivered to the surface through the V_f_lander trips, each of which consumes M_p_lander in its own descent from LLO.

4. Therefore, the term M_p_lander × N_V_f_lander in the denominator of Theorem 6 includes the logistical cost of transporting to the surface all the propellant required for the ascents, including that of the cargo trip itself and that of the refueling trips.
5. **Necessary condition:** For the system to be viable without ISRU, the following must hold:

$P_lander > M_{p2}$

and also:

$M_p_lander > M_{p2}$

The first condition ensures that a refueling trip can deliver more propellant than it consumes in its own ascent. The second ensures that the lander has sufficient propellant capacity to complete the descent.

In Model 1:

P_lander = 36.73 t > 11.15 t = M_{p2} → Condition satisfied. M_p_lander = 38.63 t > 11.16 t = M_{p2} → Condition satisfied.

Conclusion: M_{p2} is included in Theorem 6 through N_V_f_lander. It is not necessary to add it explicitly to the denominator, because the surface propellant balance has already fully accounted for it.

Relation between Theorems 5 and 6

Theorems 5 and 6 describe the same characteristic relation of the stationary regime for the supplying segment LEO → LLO under two different logistical hypotheses.

With ISRU = 100% on the surface (Theorem 5):

$N_{(Vp)}$ / $N_{(Vf)}$ = P_lander / M_p_lander

Without ISRU on the surface (Theorem 6):

$N_{(Vp)}$ / $N_{(Vf)}$ = P_lander / [M_p_lander × (1 + N_V_f_lander)]

The only difference between both expressions is the factor (1 + N_V_f_lander), which represents the additional logistical cost associated with transporting the propellant required to sustain the lander's refueling trips when no local propellant production exists.

When:

N_V_f_lander = 0

the expression of Theorem 6 immediately reduces to:

$N_{(Vp)}$ / $N_{(Vf)}$ = P_lander / M_p_lander

which coincides exactly with the expression obtained in Theorem 5.

Consequently, **Theorem 5 constitutes a particular case of Theorem 6**, corresponding to the regime in which ascent propellant is available on the surface without additional logistical cost (ISRU = 100%).

BUFFERS AS DAMPERS, NOT ACCUMULATORS

A fundamental aspect of the present theory, which distinguishes it from intuitive approaches, is the role of buffers (payload and propellant depots).

In the usual interpretation of logistical systems, a depot is a warehouse: it fills, it empties, and the objective is to always have enough mass to avoid interrupting operation. Greater capacity is associated with greater safety.

In the present theory, this interpretation is incorrect.

The payload and propellant depots of Depot_sub_LLO are not accumulators in the traditional sense. They are **buffers** whose function is to absorb the discrete oscillations of the system.

Their function is not to accumulate mass, but to **dampen the fluctuations** resulting from the difference between:

- The mass P_{op} of payload or inputs that arrives with each Vp trip.
- The mass P_f of propellant that is consumed to deliver Pop but arrives in V_p trips that deliver an amount of propellant P different from P_f.

Due to the discretization of events and the variability of the sequence, there is always a difference between the mass that enters and the mass that leaves. The buffer absorbs that difference.

Consequently:

• **More capacity is not better**. Greater capacity implies greater infrastructure cost and greater structural mass.

• **The optimal capacity is the minimum necessary** to absorb the oscillations of the sequence without interrupting operation.

• **The buffer is not a strategic warehouse**. It is a dynamic damper.

• The accumulation of propellant and/or payload in the Depots is not an indicator of network strength. On the contrary, such accumulation is the clearest indicator that **the network is being under-utilized**.

• **The network must always operate under tension**, never accumulate payloads in transit nor propellant beyond what is necessary to absorb oscillations. More accumulated propellant is not better.

This distinction is fundamental for sizing depots in a continuously operating logistical network.

IMPROVEMENT FACTOR OF THE REGIME INTRODUCED BY ISRU

Statement

The multiplicative factor introduced by the implementation of ISRU on the characteristic relation of the stationary regime is given by:

$$F_{ISRU} = 1 + N_V_f_lander$$

where N_V_f_lander is the number of lander refueling trips required per cargo trip in the regime without ISRU (defined in Theorem 4).

Equivalently, as a function of the fundamental parameters of the lander:

$$F_{ISRU} = P_lander / (P_lander - M_{p2})$$

where:

- P_lander is the mass the lander delivers to the surface per descent trip (either net payload or propellant).
- M_{p2} is the propellant the lander needs to ascend empty from the surface to LLO.

Proof

1. Starting from the definition of the improvement factor

The improvement factor F_{ISRU} is defined as the quotient between the characteristic relation of the stationary regime with ISRU (Theorem 5) and the characteristic relation of the stationary regime without ISRU (Theorem 6):

$$F_{ISRU} = [N_{(Vp)}/N_{(Vf)}]_WITH_ISRU / [N_{(Vp)}/N_{(Vf)}]_WITHOUT_ISRU$$

2. Substitute the expressions from Theorems 5 and 6

From Theorem 5 (WITH ISRU):

$$[N_{(Vp)}/N_{(Vf)}]_WITH_ISRU = P_lander / M_p_lander$$

From Theorem 6 (WITHOUT ISRU):

$$[N_{(Vp)}/N_{(Vf)}]_WITHOUT_ISRU = P_lander / [M_p_lander \times (1 + N_V_f_lander)]$$

3. Form the quotient

$$F_{ISRU} = [P_lander / M_p_lander] / [P_lander / (M_p_lander \times (1 + N_V_f_lander))]$$

4. Simplify

$$F_{ISRU} = (P_lander / M_p_lander) \times (M_p_lander \times (1 + N_V_f_lander) / P_lander)$$

$$F_{ISRU} = 1 + N_V_f_lander$$

5. Explicit expression as a function of P_lander and Mp2

From Theorem 4:

$$N_V_f_lander = M_{p2} / (P_lander - M_{p2})$$

Substituting:

$$F_{ISRU} = 1 + M_{p2} / (P_lander - M_{p2})$$

Sum:

$$F_{ISRU} = (P_lander - M_{p2} + M_{p2}) / (P_lander - M_{p2})$$

$$F_{ISRU} = P_lander / (P_lander - M_{p2})$$

Q.E.D.

Observations

1. Interpretation of the improvement factor

The factor F_{ISRU} quantifies the multiplier that the implementation of ISRU applies to the efficiency of the system. That is:

$$\mathbf{[\ N_{(Vp)}/N_{(Vf)}\ ]_{WITH_ISRU} = F_{ISRU} \times [\ N_{(Vp)}/N_{(Vf)}\ ]_{WITHOUT_ISRU}}$$

Or, equivalently, the factor F_{ISRU} quantifies the increase in the characteristic relation of the regime introduced by local propellant production. Consequently:

$$\mathbf{[N_{(Vf)}/N_{(Vp)}]_{ISRU} = F_{ISRU} \times [N_{(Vf)}/N_{(Vp)}]_{WITHOUT_ISRU}}$$

That is, the implementation of ISRU increases by a factor F_{ISRU} the proportion of cargo trips relative to refueling trips that the supplying segment can sustain.

2. Equivalent expression

The factor F_{ISRU} can be expressed in two equivalent forms:

$$\mathbf{F_{ISRU} = 1 + N_V_f_lander = P_lander / (\ P_lander - M_{p2}\ )}$$

The first form shows its origin in logistical recursion (Theorem 4). The second expresses it directly as a function of the fundamental parameters of the lander, which are more accessible and easier to obtain from vehicle design data.

3. Exclusive dependence on the lander

The improvement factor F_{ISRU} depends exclusively on the properties of the lander:

- P_lander: what the lander delivers to the surface per trip.
- Mp2: what the lander needs to ascend empty.

It does not depend on the supplying segment (tug), nor on the number of annual launches, nor on the transport capacity P_{tug}. This reflects that the improvement introduced by ISRU occurs in the served segment (LLO → SUP) and propagates upward to the supplying segment (LEO → LLO), but its magnitude is determined exclusively by the vehicle operating at the end of the network.

4. Condition of validity

The improvement factor is valid if and only if:

$P_lander > M_{p2}$

If $P_lander \leq M_{p2}$, the denominator $P_lander - M_{p2}$ is zero or negative. In that case, the system without ISRU cannot sustain itself in stationary regime (as established in Theorem 4), and the improvement factor has no physical meaning.

5. Application to Model 1

For Model 1:

• P_lander = 36.73 t

• M_{p2} = 11.16 t

F_{ISRU} = 36.73 / (36.73 − 11.15) = 36.73 / 25.58 = 1.436

That is, the system with ISRU is 1.436 times more efficient than the system without ISRU in terms of the relation $N_{(Vp)}/N_{(Vf)}$.

Or, equivalently:

$$\mathbf{N_{(Vp)}/N_{(Vf)}\ with\ ISRU = 1.436 \times N_{(Vp)}/N_{(Vf)}\ without\ ISRU}$$

THEOREM 7 — Existence and Uniqueness of the Steady-State Regime with Finite Buffers

Introduction

Theorems 5 and 6 establish the ratio $N_{(Vp)}/N_{(Vf)}$ as the mean value toward which the steady-state regime of the network tends. However, that formulation leaves an open question: why should the system converge to that value? Nothing in the average mass balance, by itself, prevents the network from operating with a different ratio, even in a sustained way.

The present theorem closes that question. It demonstrates that the ratio $N_{(Vp)}/N_{(Vf)}$ is not simply the most convenient or efficient value, but the only value compatible with the physical fact that no real depot has infinite capacity. Any launch sequence that maintains a different ratio, in a sustained manner, inevitably ends up emptying or overflowing one of the $Depot_{LLO}$ buffers in a finite number of steps. In this sense, the attractor of Theorems 5 and 6 ceases to be an observed tendency and becomes a necessary condition for the survival of the regime.

STATEMENT

Let $\mathbf{n = 1, 2, 3, ...}$ be the discrete index of events (ordering of launches of the supplying segment, not continuous time). Let $\mathbf{L_{(n)} \in \{V_p, V_f\}}$ be the sequence of trip decisions, and let $\mathbf{S_{(n)}}$ be the number of $\mathbf{V_p}$ trips among the first $\mathbf{n}$ trips, so that $\mathbf{N_{(Vf,n)} = n - S_{(n)}}$.

A sequence is said to be admissible if, for all n, the state of the $Depot_{LLO}$ buffers satisfies:

$$\mathbf{0 \le Bp_{(n)} \le Cap_p}$$

$$\mathbf{0 \le Bf_{(n)} \le Cap_f}$$

where $\mathbf{Bp_{(n)}}$ and $\mathbf{Bf_{(n)}}$ are the masses accumulated in the payload buffer and the propellant buffer respectively, defined as:

$Bp_{(n)} = Bp_{(0)} + [V_p$ trips up to n] × P_{tug} − [lander descents up to n] × P_lander

$Bf_{(n)} = Bf_{(0)} + [V_f$ trips up to n] × P_{tug} − [lander descents up to n] × M_p_lander

That is: a sequence is admissible if it never allows a buffer to run out of mass (0) nor to exceed its physical capacity (Cap).

Then, every admissible sequence necessarily satisfies:

lim (n→∞) $S_{(n)}$/n = ρ*,

with ρ* = P_lander / (P_lander + M_p_lander)

That is: **$N_{(Vp)}/N_{(Vf)}$ → P_lander/M_p_lander** (Theorem 5) is not an observed tendency of the system, but a necessary condition for the buffers to remain bounded indefinitely.

Proof

1. Since **0 ≤ $Bf_{(n)}$ ≤ Cap_f** for all n, the contribution of **B_f** to the accumulated balance is bounded by a constant independent of **n**. Consequently:

$N_{(Vf,n)}$ × P_{tug} = N(V_p_lander,n) × M_p_lander + $O_{(1)}$

where the term $O_{(1)}$ captures the bounded difference between what is delivered and what is consumed, absorbed by the buffer.

2. Using N(V_p_lander,n) = [$S_{(n)}$ × P_{tug}] / P_lander (relation (3) of Theorem 5):

(n − S(n)) × P_{tug} = [S(n) × P_{tug} / P_lander] × M_p_lander + $O_{(1)}$

3. Dividing both sides by n and taking n → ∞, the term $O_{(1)}$/n → 0. If ρ = lim $S_{(n)}$/n exists:

(1 − ρ) = ρ × M_p_lander / P_lander

4. Solving:

$$\rho = P_lander / (P_lander + M_p_lander) = \rho^*$$

Q.E.D.

VALIDITY CONDITIONS

Theorem 7 is valid under the following conditions:

1. The $Depot_{LLO}$ buffers have finite capacity (Capp and Capf are finite and known). This is the central hypothesis of the theorem: without finite capacity, nothing prevents the network from operating indefinitely with a ratio different from ρ*.
2. The buffer can never take negative values nor exceed its physical capacity; both violations are interpreted as collapse of the regime (lack of propellant or lack of storage space).

3. The conversion relations between trips of the supplying segment and lander descents (P_{tug}, P_lander, M_p_lander) are those established in Theorem 5, and remain constant in the interval considered.
4. The limit ρ = lim S(n)/n is assumed to exist; the theorem proves what value it must take if it exists, not that it necessarily exists for every arbitrary sequence (see Corollary 7.1 for the case where ρ ≠ ρ*).

OBSERVATIONS

1. From tendency to necessity

This is the central contribution of the theorem: it converts the attractor of Theorems 5 and 6, which until now could be read as an observed average of the system, into a necessary condition derived exclusively from the fact that the depots are finite. No additional hypothesis about system behavior—no control rule, no correction mechanism—is required to reach this result. It is enough that the buffers cannot grow or empty without limit.

2. Why the payload buffer (B_p) does not appear in the proof

The proof relies on the boundedness of B_f, the propellant buffer. B_p is bounded by the same argument, symmetrically, but it is not necessary to invoke it separately: once the ratio ρ* is fixed through the propellant balance, the net payload balance becomes automatically consistent, because both flows share the same P_{tug} at the origin.

COROLLARY 7.1 — Collapse of Non-Admissible Sequences

If a sequence maintains an asymptotic ratio ρ ≠ ρ*, *then* $Bf_{(n)}$ *(or* $Bp_{(n)}$*) exhibits a linear drift in n of magnitude proportional to* |ρ – ρ*|, and therefore violates the bound Capf (or Capp) in a finite n. Consequently, no sequence with ρ ≠ ρ* is admissible indefinitely: the system collapses in finite time.

In practical terms: there is no network architecture that can “choose” to operate with a trip ratio different from ρ* and sustain it over time. Any sustained deviation eventually results in an empty or overflowing depot.

COROLLARY 7.2 — The Basin of Attraction as the Set of Admissible Sequences

The basin of attraction of the regime is not a diffuse neighborhood around a mean value, but the precise set of sequences $\{L_{(n)}\}$ such that $Bp_{(n)}$ and $Bf_{(n)}$ remain within their bounds for all n. Geometrically, in the plane ($N_{(Vp)}$, $N_{(Vf)}$), these sequences remain around the line through the origin with slope $N_{(Vp)}/N_{(Vf)}$ = ρ*/(1−ρ)* = P_lander / M_p_lander; moving away from that line in a sustained manner implies leaving the basin of attraction and losing admissibility.

Fundamental Consequence of Theorem 7 + Corollaries 7.1 and 7.2

It is demonstrated that:

1. The characteristic relation of the steady-state regime (Theorems 5 and 6) is not a mere statistical tendency of the system.
2. It is the only relation compatible with the existence of finite buffers.
3. Any sustained deviation from that relation introduces an accumulated drift in the buffers.
4. Any accumulated drift inevitably leads to emptying or overflowing one of the buffers in a finite number of operations.
5. Consequently, the basin of attraction of the regime coincides exactly with the set of admissible sequences.

Therefore, the steady-state regime is not a design option nor an operational strategy: it is an inevitable consequence of the physical constraints imposed by the existence of finite buffers.

THEOREM 8 — Existence and Uniqueness of the Steady-State Regime with Logistic Recursion (without ISRU)

Statement

In the absence of ISRU, the lander operating between LLO and the surface introduces logistic recursion: for each payload trip it must perform additional refueling trips to deliver to the surface the propellant required for its own ascent. This recursion modifies the characteristic relation of the steady-state regime of the LEO → LLO segment.

Let **n = 1, 2, 3, ...** be the discrete index of events (ordering of launches of the supplying segment, not continuous time). Let $\mathbf{S_{(n)}}$ be the number of $\mathbf{V_p}$ trips among the first n trips, so that $\mathbf{N_{(Vf,n)} = n - S_{(n)}}$.

A sequence is admissible if the Depot_{LLO} buffers satisfy, for all n:

$$\mathbf{0 \le Bp_{(n)} \le Cap_p}$$

$$\mathbf{0 \le Bf_{(n)} \le Cap_f}$$

where $Bp_{(n)}$ and $Bf_{(n)}$ are the masses accumulated in the payload buffer and the propellant buffer respectively, defined as:

$$\mathbf{Bp_{(n)} = Bp_{(0)} + [V_p \text{ trips up to n}] \times P_{tug} - [\text{lander descents up to n}] \times P_lander}$$

$$\mathbf{Bf_{(n)} = Bf_{(0)} + [V_f \text{ trips up to n}] \times P_{tug} - [\text{lander descents up to n}] \times M_p_lander \times (1 + N_V_f_lander)}$$

Then, every admissible sequence necessarily satisfies:

$$\mathbf{\lim (n \to \infty)\ S_{(n)}/n = \rho^*}$$

where:

$\rho^* = P_lander / (P_lander + M_p_lander \times (1 + N_V_f_lander))$

That is: the fraction of payload trips of the supplying segment necessarily tends to the value ρ*, which is the efficiency of the served segment without ISRU (Theorem 4). Any different sustained ratio causes the buffers to collapse in finite time.

Proof

1. Since $\mathbf{0 \le Bf_{(n)} \le Cap_f}$ for all **n**, the difference between propellant delivered and propellant consumed is bounded by a constant independent of **n**.

Therefore:

$$N_{(Vf,n)} \times P_{tug} = N(V_p_lander,n) \times M_p_lander \times (1 + N_V_f_lander) + O_{(1)}$$

where the term $O_{(1)}$ captures the bounded difference between what is delivered and what is consumed, absorbed by the buffer.

2. Using the relation between tug trips and lander descents:

$$N(V_p_lander,n) = [S_{(n)} \times P_{tug}] / P_lander$$

we substitute:

$$(n - S_{(n)}) \times P_{tug} = [S_{(n)} \times P_{tug} / P_lander] \times M_p_lander \times (1 + N_V_f_lander) + O_{(1)}$$

3. Dividing both sides by n and taking $n \to \infty$, the term $O_{(1)}/n \to 0$. If $\rho = \lim S(n)/n$ exists, then:

$$1 - \rho = \rho \times [M_p_lander \times (1 + N_V_f_lander)] / P_lander$$

4. Solving for ρ:

$$\rho = P_lander / [P_lander + M_p_lander \times (1 + N_V_f_lander)]$$

Q.E.D.

VALIDITY CONDITIONS

1. The $Depot_{LLO}$ buffers have finite capacity.
2. The buffers cannot take negative values nor exceed their physical capacity.
3. The relations P_{tug}, P_lander and M_p_lander remain constant.
4. The limit $\rho = \lim S(n)/n$ is assumed to exist.

COROLLARY 8.1 — Collapse of Non-Admissible Sequences

If a sequence maintains an asymptotic ratio **ρ ≠ ρ***, *then* ***Bf*$_{(n)}$** *or* ***Bp*$_{(n)}$** *exhibits a linear drift in* ***n****, proportional to* **|ρ – ρ*|.** Since the buffers are finite, this drift reaches a bound in a finite number of steps. Therefore, no sequence with **ρ ≠ ρ*** can be sustained indefinitely.

COROLLARY 8.2 — Basin of Attraction of the Regime without ISRU

The basin of attraction consists of all sequences **{$L_{(n)}$}** such that **$Bp_{(n)}$** and **$Bf_{(n)}$** remain within their bounds for all **n**. In the plane (**$N_{(Vp)}$, $N_{(Vf)}$**), these sequences remain around the line:

$N_{(Vp)}/N_{(Vf)}$ = P_lander / [M_p_lander × (1 + N_V_f_lander)]

Moving away from this line in a sustained manner implies leaving the basin of attraction and collapsing the regime.

GLOBAL READING OF THEOREM 8

Theorem 8 is the complete version of Theorem 7 for the case without ISRU. It demonstrates that:

1. The lander recursion propagates upward through the network.
2. The steady-state regime of the supplying segment is necessarily modified.
3. The characteristic ratio of the regime is the only value that avoids collapse.
4. Any sustained deviation produces linear drift and collapse in finite time.
5. The basin of attraction is the exact set of sequences that keep the buffers bounded.

In consequence:

Logistic recursion does not modify the nature of the steady-state regime, only its characteristic value. The existence of finite buffers still imposes a single admissible regime; the lander recursion shifts that regime by incorporating the additional logistic cost associated with transporting ascent propellant.

UNIFICATION WITH THEOREM 7 (limit case ISRU = 100%)

Theorems 7 and 8 are not independent results, but two particular cases of the same mathematical principle: the existence and uniqueness of the steady-state regime imposed by the boundedness of the $Depot_{LLO}$ buffers.

This is verified by taking the limit N_V_f_lander → 0 in Theorem 8:

ρ* = P_lander / [P_lander + M_p_lander × (1 + N_V_f_lander)]

When $N_V_f_lander \to 0$ (absence of recursion, ascent propellant available without logistic cost — the ISRU = 100% case already identified in Corollary 1 of Theorem 6):

$\rho^* \to$ P_lander / (P_lander + M_p_lander)

which is exactly the result of Theorem 7.

Therefore, Theorem 7 is not a different case from Theorem 8, but its particular instance: both are manifestations of a single principle of necessity, the boundedness of the $Depot_{LLO}$ buffers forcing a unique equilibrium ratio, applied to the propellant demand of the served segment, whether that demand is simple (M_p_lander, with ISRU) or recursive (M_p_lander × (1+N_Vf_lander), without ISRU).

RELATION BETWEEN ρ AND THE CHARACTERISTIC RATIO OF THE REGIME

Throughout Theorems 5 to 8, two different forms have been used to describe the same steady-state regime.

The first is the characteristic ratio of the regime:

$N_{(Vp)} / N_{(Vf)}$

which expresses the relation between trips dedicated to transporting payload and trips dedicated to transporting operational resources.

The second is the characteristic fraction of the regime:

$\rho^* = N_{(Vp)} / [N_{(Vp)} + N_{(Vf)}]$

which represents the fraction of all trips of the supplying segment that correspond to payload trips.

Both expressions describe exactly the same physical state of the system and contain the same information. The difference lies only in the way the regime is represented.

The relation between the two is obtained directly from the definition of ρ*.

Starting from:

$\rho^* = N_{(Vp)} / [N_{(Vp)} + N_{(Vf)}]$

dividing numerator and denominator by $N_{(Vf)}$:

$\rho^* = [N_{(Vp)}/N_{(Vf)}] / [1 + N_{(Vp)}/N_{(Vf)}]$

Solving for the ratio:

$N_{(Vp)} / N_{(Vf)} = \rho^* / (1 - \rho^*)$

Equivalently:

$\rho^* = [N_{(Vp)}/N_{(Vf)}] / [1 + N_{(Vp)}/N_{(Vf)}]$

Consequently, both expressions are completely equivalent and either may be used depending on convenience.

The ratio $N_{(Vp)}/N_{(Vf)}$ is particularly suitable for studying logistic balances and the distribution of payload and operational resource flows.

By contrast, the fraction ρ* is more natural in the mathematical analysis of discrete sequences and in the demonstration of the existence and uniqueness of the steady-state regime, since it directly represents the asymptotic fraction of payload trips within the total set of operations.

Therefore, $N_{(Vp)}/N_{(Vf)}$ and ρ* do not represent two different properties of the system, but two equivalent manifestations of the same steady-state logistic regime. The first describes the relation between the two types of trips; the second expresses what fraction of the total operations corresponds to payload trips.

THE ANNUAL OPERATIVE VECTOR

Introduction

Theorem 7 establishes that, in steady-state regime, the ratio $N_{(Vp)}/N_{(Vf)}$ necessarily converges to its characteristic value. However, that ratio describes only the asymptotic structure of the regime; it is not yet the object that must be programmed in the actual operation of the network. In a simulation or in an annual launch plan, one does not program a ratio, but a concrete set of operations: how many payload trips and how many refueling trips must be performed during each operational cycle.

To describe that operative state, the **annual operative vector $V_{(t)}$** is introduced, defined below. As will be demonstrated in Theorem 9, the name is not merely descriptive: once the network reaches its steady-state regime, $V_{(t)}$ behaves exactly like a vector in the strict mathematical sense — a point in the plane $(N_{(Vp)}, N_{(Vf)})$ with a fixed direction (determined by the physics of the system) and a single free magnitude ($L_{(t)}$).

Once the steady-state regime is reached, this vector no longer possesses two independent degrees of freedom. Its direction is fixed by the physics of the system, represented by the characteristic ratio of the regime, and only its magnitude remains free, determined by the desired level of logistic activity.

Definition — Vector $V_{(t)}$

Let t = 1, 2, 3, ... be the discrete index of the operational year (launch season), distinct from the event index n used in Theorem 7. The operative state of the network in year t is described by the vector:

$$\mathbf{V(t) = (N_{(Vp,t)}, N_{(Vf,t)})}$$

that is, the number of payload trips and refueling trips programmed for the supplying segment during that year. This vector is the object actually programmed in a simulation

or launch plan, in contrast to the asymptotic ratio $N_{(Vp)}/N_{(Vf)}$, which describes long-term behavior.

THEOREM 9 — Reduction to One Degree of Freedom in Steady-State Regime

Statement

Once the network has reached its steady-state regime (the basin of attraction of Theorem 7), the vector $V_{(t)}$ ceases to have two independent components. It is determined by a single operative scalar:

$$\mathbf{L_{(t)} = N_{(Vp,t)} + N_{(Vf,t)}}$$

the total number of launches of the supplying segment in year t, through:

$$\mathbf{N_{(Vp,t)} = L_{(t)} \times \rho^*}$$

$$\mathbf{N_{(Vf,t)} = L_{(t)} \times (1 - \rho^*)}$$

where ρ* = P_lander / (P_lander + M_p_lander) is a fixed constant, determined exclusively by the design of the lander (Theorem 5), not by the operation.

Proof

1. By definition,
   $\mathbf{L_{(t)} = N_{(Vp,t)} + N_{(Vf,t)}}$.
2. By Theorem 7, in steady-state every admissible sequence satisfies
   $\mathbf{N_{(Vp,t)}/N_{(Vf,t)} = \rho^*/(1-\rho^*)}$.
3. From (1) and (2):
   $\mathbf{N_{(Vp,t)} = \rho^* \times L_{(t)}}$
   and
   $\mathbf{N_{(Vf,t)} = (1-\rho^*) \times L_{(t)}}$.

Q.E.D.

Theorem 9 is valid under the following conditions:

1. The network has already reached its steady-state regime, i.e., it lies within the basin of attraction defined in Corollary 7.2. The theorem does not describe the bootstrap or initial deployment period, where $V_{(t)}$ may deviate from this reduction (see later the bootstrap/steady-state decomposition).

2. The constants P_lander and M_p_lander that determine ρ* remain fixed in the interval considered. If the lander changes model or configuration, ρ* changes with it, and $V_{(t)}$ realigns to the new direction.
3. $N_{(Vp,t)}$ and $N_{(Vf,t)}$ are integer values in real operation; the expression $N_{(Vp,t)} = L_{(t)} \times \rho^*$ must be understood as the mean value toward which annual programming approaches, not as an exact event-by-event equality (same discrete character noted in the Observations of Theorem 5).

OBSERVATIONS

1. $\mathbf{V_{(t)}}$ as a vector in the strict sense

The name “vector” is not a rhetorical figure. Once ρ* is fixed, $V_{(t)}$ is literally a point in the plane ($N_{(Vp)}$, $N_{(Vf)}$) lying on the line through the origin with slope $\rho^*/(1-\rho^*)$ — the same line described in Corollary 7.2. Within that line, $L_{(t)}$ is the only free coordinate: it is the scalar that multiplies the unit vector of fixed direction ($\rho^*, 1-\rho^*$).
In this sense, $V_{(t)} = L_{(t)} \times (\rho^*, 1-\rho^*)$: a vector of constant direction and variable magnitude.

2. What it means, operationally, to have only one degree of freedom

Before the steady-state regime is established, one might suppose that programming the network requires deciding two quantities per year: how much payload to send and how much propellant to send.
Theorem 8 shows that this apparent freedom of two variables is illusory: once the desired payload delivery is fixed (or, equivalently, once the total number of launches is fixed), the required amount of propellant is determined; it is not an additional decision.

COROLLARY 9.1 — Delivered Payload as a Linear Function of $L_{(t)}$

The payload delivered in year t is:

$$\mathbf{Payload_{(t)} = N_{(Vp,t)} \times P_{tug} = L_{(t)} \times \rho^* \times P_{tug}}$$

That is, the delivery capacity of the network scales linearly with the total number of annual launches, with fixed slope $\rho^* \times P_{tug}$ determined by the vehicle design, not by the operation.

COROLLARY 9.2 — $L_{(t)}$ as the Only Control Variable

Consequently, once in steady-state regime, there is no combination of $N_{(Vp,t)}$ and $N_{(Vf,t)}$ that can be adjusted independently: both are the projection of a single scalar $L_{(t)}$ onto the fixed direction ρ* of the attractor line (Corollary 7.2). The only way to increase or decrease the payload flow delivered in a year is to increase or decrease the total number of launches $L_{(t)}$; there is no other operational freedom in this sense.

Final Observation

The internal variables of the network (number of lander descents, number of lander refueling trips in the non-ISRU case, etc.) are also functions derived from $N_{(Vp,t)}$ through the relations already established in Theorems 5 and 6, and therefore functions derived, ultimately, from the same scalar $L_{(t)}$. The annual operative vector $V_{(t)}$, in mature steady-state regime, has only one real degree of freedom.

GENERALIZATION OF THE OPERATIVE VECTOR

The vector defined in this chapter constitutes the minimal representation of the steady-state regime of the supplying segment. However, a complete logistic network includes other operations that can also be expressed as discrete quantities per operational cycle.

Consequently, it is natural to define an expanded operative vector, whose components represent the annual number of operations performed by each type of vehicle or segment of the network.

In its general form:

$$\mathbf{Vop_{(t)} = (N_{(Vp)}, N_{(Vf)}, N(_{Lp)}, N_{(Lf)}, N_{(Lup)}, N_{(Lret)}, N_{(Tret)}, \dots)}$$

where:

- $\mathbf{N_{(Vp)}}$ represents payload trips of the supplying segment.
- $\mathbf{N_{(Vf)}}$ represents refueling trips of the supplying segment.
- $\mathbf{N_{(Lp)}}$ represents lander descents carrying payload.
- $\mathbf{N_{(Lf)}}$ represents lander descents carrying propellant.
- $\mathbf{N_{(Lup)}}$ represents lander ascents from the surface.
- $\mathbf{N_{(Lret)}}$ represents lander returns carrying payload.
- $\mathbf{N_{(Tret)}}$ represents tug returns or other reusable vehicle returns.

Each component represents the number of operations performed during a complete cycle of the regime.

In a specific architecture, additional components may exist or some of the above may disappear. The dimension of the vector will depend on the complexity of the network considered.

What matters is that all these components do not constitute independent variables. Once the steady-state regime is established, they are determined by the structural relations of the network and, ultimately, by the level of logistic activity represented by $L_{(t)}$.

Thus, the operative vector constitutes a complete description of the annual operational state of the network and summarizes, through a single mathematical structure, all operations performed during a cycle.

THEOREM 10 — Minimum Buffer Sizing

Statement

Let a buffer in a segmented logistics network receive and deliver mass through an admissible discrete sequence of operations belonging to the steady-state regime defined in Theorems 7 and 8.

Let $B_{(n)}$ be the mass stored in that buffer after the n-th event of the sequence.

Then, the minimum buffer capacity required to guarantee indefinite operation of the regime is:

$$\mathbf{C_{min} = max(B_{(n)}) - min(B_{(n)}) + \delta}$$

where:

• $\max(B_{(n)})$ is the maximum level reached by the buffer during the sequence considered. • $\min(B_{(n)})$ is the minimum level reached. • $\delta > 0$ is a safety margin chosen by the designer to ensure that the buffer never reaches complete depletion during operation.

Consequently, the optimal capacity of a buffer does not depend on the total mass circulating through the network, but solely on the amplitude of the oscillations produced by the discretization of flows.

Proof

By definition, the buffer content after event n is:

$$\mathbf{B_{(n)} = B_{(0)} + \Sigma\, Min_{(k)} - \Sigma\, Mout_{(k)}}$$

where $\Sigma\, Min_{(k)}$ and $\Sigma\, Mout_{(k)}$ represent, respectively, the accumulated mass introduced into and withdrawn from the buffer up to event n.

During steady-state regime, the buffer content remains bounded and oscillates between a minimum and a maximum value:

$$\mathbf{\min(B_{(n)}) \leq B_{(n)} \leq \max(B_{(n)})}$$

To maintain indefinite operation, two conditions must hold simultaneously:

• the buffer must never deplete; • the buffer must never exceed its physical capacity.

The maximum excursion experienced by the buffer content during the sequence is, by definition:

$$\mathbf{\Delta B = \max(B_{(n)}) - \min(B_{(n)})}$$

Therefore, any capacity smaller than ΔB would necessarily cause depletion or overflow of the buffer at some event in the sequence.

Adding a positive operational margin δ, intended to absorb operational uncertainties or design criteria, yields the minimum required capacity:

$$\mathbf{C_{min} = \Delta B + \delta}$$

that is,

$$\mathbf{C_{min} = \max(B_{(n)}) - \min(B_{(n)}) + \delta}$$

Q.E.D.

VALIDITY CONDITIONS

- The sequence considered belongs to the basin of attraction of the steady-state regime.
- Buffers have finite capacity.
- Buffer content cannot take negative values.
- The masses exchanged in each operation are defined by the architecture considered.
- The margin δ is a design parameter, not a property of the logistic regime.

OBSERVATIONS

1. The buffer as a damper The theorem shows that buffer sizing depends exclusively on the amplitude of the oscillations it must absorb, not on the total mass flowing through the network during an operational cycle.
2. Optimal capacity A capacity greater than Cmin does not improve the logistic regime. It only increases structural mass, cost, and infrastructure complexity without providing any improvement in regime stability.
3. Independence from architecture The expression obtained is completely general. It does not depend on vehicle type, propellant type, or network geometry. It only requires knowledge of the temporal evolution of buffer content.

RETURN COROLLARY — Marginal Perturbation on the Unidirectional Regime

The steady-state regime characterized by Theorems 7, 8, and 9 describes a strictly unidirectional network: every tug is discarded after completing its Earth → Moon transport function, and every lander ascent is an empty repositioning. The parameters that characterize each model (ρ^*, $N_{(Vp)}$, $N_{(Vf)}$) are computed over this base regime.

The return of lunar payload to Earth does not constitute a different regime nor does it require recalculating ρ^*. It is treated as a discrete and optional perturbation on the base regime: one of the tugs that, in the unidirectional regime, would be discarded after

completing a Vp or V_f trip, is retained in LLO and refueled with the propellant required to place a mass Pexport on an Earth-return trajectory. In the same step, one of the lander's empty ascents becomes a loaded ascent carrying lunar payload, feeding the lunar-payload buffer from which the export is drawn.

Applying Tsiolkovsky to the tug's reentry-injection maneuver:

$Mp_{ret} = (Ms_{tug} + P_{export}) \times (R_{ret} - 1)$, with $R_{ret} = e^{(\Delta v_{ret} / (Isp \times g_0))}$

For Model 1, with $Ms_{tug} \approx 4.22$ t, Isp = 340 s, $\Delta v_{ret} = 900$ m/s, and $P_{export} = 15$ t (internal composition: 10 t net payload plus a reentry capsule of approximately 5 t, no life support, comparable in scale to a cargo-type Dragon capsule):

$Mp_{ret} \approx 5.96$ t

The internal composition of P_{export} (how much is capsule, how much is net payload) is a Model/Implementation decision, not a theoretical constraint. From the standpoint of the present theory, both form a single exported mass. Tsiolkovsky depends only on the total mass moved, not on its internal distribution.

The marginal cost of each export, measured as a fraction of a V_f trip of the supplying segment ($P_{tug} = 43.00$ t):

$\Delta N_{(Vf)} = Mp_{ret} / P_{tug} \approx 5.96 / 43.00 \approx 13.9\%$

That is, each 15-t export to Earth costs approximately 14% of a V_f trip — equivalent to sacrificing, on average, one full V_f trip for every 7 exports performed.

Operationally, this is executed by modifying the annual vector at a single point: instead of the nominal allocation that satisfies ρ*, some years the vector deviates slightly to accommodate the cost of an export. This punctual deviation does not remove the system from the basin of attraction defined in Corollary 7.2/8.2 — the regime remains stable — as long as the return is kept at the marginal scale described here and is not systematized every cycle.

It should be noted that Principle 3 (flow decoupling, tug discard after its function) admits this punctual exception: not all tugs are discarded; a marginal fraction is retained to serve the return, without invalidating flow decoupling as the general regime of the supplying segment.

FINAL OBSERVATION — Contrast with the Current State of the Art

The cost calculated here — approximately 14% of an already-existing Vf trip for each 15 t exported — contrasts with any sample-return mission conceived as an independent operation: a dedicated mission, with its own launch, its own trajectory window, and typically designed to return a payload mass on the order of tens of kilograms, not tons.

No lunar-exploration architecture currently under development contemplates a mechanism for returning payload at this scale as an almost free byproduct of regular logistic operation. **The difference does not lie in the technology employed, but in**

the object of study. Whereas traditional architectures treat return as an independent mission, the present theory interprets it as a marginal perturbation on a permanent logistic regime. Under this approach, return ceases to be a specific architecture and becomes an additional operation integrated into an already existing logistics network.

Such a mechanism only emerges naturally when the architecture is conceived from the outset as a network of decoupled flows with shared tugs and buffers (Principle 3), and not as a succession of independent missions.

BIDIRECTIONALITY: COUPLED NETWORKS UNDER IDEAL SYMMETRY

INTRODUCTION

All theory developed so far describes a unidirectional network: mass flowing from Earth to the Moon, with a single attractor ρ* governing the split between payload trips and refueling trips. This section explores what happens when that restriction is lifted and a return flow — lunar payload back to Earth — is admitted under a specific boundary condition: lunar ISRU production so advanced and self-sufficient that propellant on the Moon ceases to be a scarce resource, exactly as terrestrial propellant has been implicitly assumed throughout the preceding theory.

Under that condition, the Moon ceases to be only a destination and becomes also an origin, with the same non-scarcity property of propellant that Earth/LEO has in the rest of the theory. It follows that the return flow is not a correction or appendix to the outbound network: it is a second complete unidirectional network, with its own flow structure, sharing physical infrastructure (landers, tugs, $\text{Depot}_{\text{LLO}}$) with the first.

Definition — Network A and Network B

Network A is defined as the unidirectional Earth → Moon network already developed in Theorems 5, 6, and 7 (supplying segment LEO→LLO via tug, served segment LLO→SURFACE via lander).

Network B is defined as the unidirectional Moon → Earth network, constructed as a mirror: the same lander ascends from SURFACE to LLO carrying lunar payload (instead of descending carrying terrestrial payload), and a tug — not discarded after fulfilling its function in Network A, but retained in LLO — is refueled with the exact propellant needed to inject that payload onto an Earth-return trajectory.

Under symmetric lunar ISRU, both the empty ascent of the lander (needed to restart Network A's cycle) and the loaded ascent (belonging to Network B) consume propellant produced locally at SURFACE, with no logistic transport cost. Consequently, neither network requires the recursion of Theorem 4: both collapse, at the lander level, to the form of Theorem 3.

Balance of Network B

The supplying segment of Network B (LLO → Earth) has the same structure as that of Network A, with two types of lander ascent to LLO:

- **Vp_B**: lander ascents carrying lunar payload, feeding the PLLO_Luna buffer.
- **Vf_B**: lander ascents carrying propellant, feeding the FLLO buffer used to refuel the export tug.

The export tug consumes Mp_{ret} of FLLO per return trip. Applying the same balance method of Theorem 5:

$$\mathbf{N(Vp_B) \times Mp_{ret} = N(Vf_B) \times P_lander}$$

from which:

$$\mathbf{N(Vp_B) / N(Vf_B) = P_lander / Mp_{ret} = \rho^*_B}$$

This is the characteristic ratio of Network B, with the same algebraic form as ρ^*_A = P_lander / M_p_lander of Network A, but governed by a different cost: not the cost of descending with payload, but the cost of refueling the tug for reentry.

The coupling constraint

The lander is a physical resource shared between both networks: every descent to SURFACE (belonging to Network A) requires, sooner or later, an ascent back to LLO — and that ascent serves either Network A (empty repositioning) or Network B (Vp_B or Vf_B). Since Network A, under symmetric ISRU, has no recursion requiring empty ascents, all available ascents may be dedicated to Network B. The resulting topological constraint is:

$$\mathbf{N(Vp_A) = N(Vp_B) + N(Vf_B)}$$

Substituting the ratio of Network B:

$$\mathbf{N(Vp_A) = N(Vp_B) \times (1 + Mp_{ret} / P_lander)}$$

This equation shows that the export capacity of Network B ($N(Vp_B)$) is determined by the descent capacity of Network A ($N(Vp_A)$), discounted by a factor (1 + Mp_{ret} / P_lander) representing the cost, in lander trips, of sustaining tug refueling.

OBSERVATION — Two attractors, not one

Unlike the unidirectional regime of Theorem 7, where a single ρ* governs the entire system, the coupled Network A / Network B system is governed by two independent characteristic ratios (ρ_A and ρ_B), connected only through the allocation constraint of the shared lander. Both ratios have the same structural form (lander transport capacity over propellant cost), but are numerically distinct unless parameters coincide.

OPEN HYPOTHESIS — Oscillation between regimes

With a single network, Theorem 7 guarantees that any sequence satisfying ρ* is admissible with respect to the system's continuous buffers (FLLO, PLLO). In the coupled system, two admissibility conditions (ρ_A and ρ_B) act simultaneously on a single discrete and indivisible resource: the inventory of landers/tugs available in LLO, in addition to the continuous buffers already known. There is no a priori guarantee that satisfying one condition does not compromise the other.

There is an intuition, not yet verified, that the coupled system does not converge to a single stable equilibrium point, but oscillates between at least two regimes: one in which the vehicle inventory is drained serving predominantly Network A, and another in which it is replenished serving predominantly Network B, in a recurrent alternation.

This intuition currently lacks analytical or empirical support. It remains to be determined, in future work:

- Whether the phenomenon corresponds to a qualitative switching between discrete states (a hybrid-automaton system with transition rules), or a single orbit of a higher-dimensional continuous dynamical system (incorporating vehicle inventory as an additional state variable alongside B_p and B_f), not yet identified.
- The conditions on vehicle inventory size (large, small, variable) under which the coupled system admits a joint stability region, analogous to the basin of attraction of Corollary 7.2, but in an expanded state space.
- Empirical verification via simulation, extending the Control Vector V2 with Network B variables (export, lunar-payload ascent, tug refueling), under the hard constraint of available physical vehicle inventory.

This section is deliberately left open as a starting point for future work on bidirectionality, and should not yet be integrated into the formal body of the Earth-Moon Logistic Regime Theory.

EDITORIAL NOTE — Two horizons for bidirectionality

The development in this section corresponds to a theoretical limit case (lunar ISRU symmetric with terrestrial, return flow comparable in frequency to outbound flow) that does not represent the immediate horizon of the Earth-Moon network. In near-term practice, flow asymmetry is and will remain dominant: outbound flow governs the architecture, and return flow is marginal in comparison.

Consequently, two treatments must be distinguished, with different scopes:

1. **Immediate treatment (integrable into the formal body of the present Theory):** Return flow is treated as a marginal perturbation on the dominant unidirectional Earth → Moon network, as developed in the three-term return corollary: an additional cost, paid in reduced effective performance (lower $N_{(Vp)}/N_{(Vf)}$), extracted from resources already present in the system (excess tugs in LLO, FLLO propellant). This treatment does not require the apparatus of two

coupled networks nor the notion of ρ^*_B; it is a first-order correction on the already-demonstrated unidirectional network.

2. **Long-term treatment (future work, not integrated):** The bidirectionality section under ideal symmetry developed above, anticipating the case in which advanced lunar ISRU makes lunar propellant equivalent to terrestrial propellant and return ceases to be marginal. This treatment is, in turn, the first case study of a broader generalization: the formalization of “Unidirectional Network” as a reusable object, and of THE NETWORK (solar-system architecture, with multiple nodes — Mars, L2, LEO stations — and multiple Unidirectional Networks coupled together) as a graph of such objects. The oscillation conjecture left open in this section is likely the central question of that future generalization, not a minor detail of the Earth-Moon case.

The present Treatise addresses exclusively point 1. Point 2 is indicated as a future research program.

THEORY OF THE EARTH–MOON LOGISTIC REGIME III

MODELS

CONCEPT, MODEL AND IMPLEMENTATION: THE IMMATERIAL NATURE OF NETWORKS

To fully understand the nature of the networks proposed in this Theory, it is necessary to distinguish **precisely the levels of abstraction** in which each one operates. This separation is not a methodological detail: it is the key to identifying which elements are universal and which depend on specific circumstances.

THE CONCEPTUAL LEVEL, THE THEORY

Networks are not hardware. They are not a design. They are not a set of parts. **Networks are a concept**.

A robust concept that can be represented through models, implemented through engineering, and executed physically in the lunar environment. But each level has different properties, responsibilities, and degrees of freedom.

Up to this point we have described the networks in operational terms — nodes, connections, energy, propellants, data, mobility — but it is essential to clarify their essence: **a network is, above all, an immaterial logical structure**, a mental framework for organizing resources, flows, and capacities. An abstract architecture that exists independently of the devices that implement it.

The hardware — landers, reactors, pipes, pumps, antennas — is only a physical manifestation of the concept. One possible form among many. A contingent implementation, not the network itself.

The networks covered in this Theory possess emergent properties that do not depend on the size of the landers, the type of reactor, the manufacturer, the propellant, the design of hoses, or any specific technology. **Resilience, scalability, redundancy, flexibility, and standardization arise from the conceptual structure of the network**, not from the hardware.

A network is more than its components: **it is the relationship between them**.

MODELS: POSSIBLE REPRESENTATIONS OF THE CONCEPT

From the concept, **Models** can be constructed. A model is a concrete representation, a way of materializing the idea into a set of physical elements or operational parameters.

One model may use hypergolic landers, another cryogenic ones; one may rely on Kilopower reactors, another on modular European reactors; one may employ solar panels and batteries, another rigid pipes instead of flexible hoses. All of them are valid if they respect the boundary conditions of this Theory.

A model may reference, as an example and to facilitate the reader's "mental image" of that model, one or more hardware components, such as engines, a nuclear reactor model, or a particular launch vehicle. **This is only a reference**. An example. **The Model does not depend on any particular hardware.**

IMPLEMENTATIONS: THE DOMAIN OF ENGINEERING

Implementation is the translation of the model into real hardware. Here engineers, technicians, manufacturers, agencies, budgets, materials, tolerances, protocols, and certifications intervene. It is the level where decisions are made about the diameter of a hose, the alloy used, which pump drives the N_2O_4, which communication protocol is employed, which valves are installed, or which engine is mounted on which lander.

This level is not the objective of this Theory, nor does it correspond to my capabilities or my role. **This Theory does not intend to canonize any specific implementation**.

EXECUTION: PHYSICAL OPERATION

Execution is the real construction and operation of an implementation. Here astronauts, robots, lunar dust, wear, maintenance, schedules, budgets, real failures, and daily logistics intervene. It is the most contingent level, the most dependent on physical reality, **and the least relevant to the Theory**.

THE KEY: each level has its own emergent properties

• The oscillations and efficiency of the Logistic Network → **emerge from the concept**.

• Payload capacity and performance → **emerge from the model**.

• Costs → **emerge from the implementation**.

• Real failures and operability → **emerge from execution**.

This separation is what allows building a general theory, independent of hardware and time.

CLARIFICATION ON THE USE OF ε IN THIS THEORY

The structural fraction (ε)

The real result depends not only on Isp in the Tsiolkovsky equation, but also on the structural mass of the vehicle, and particularly on the structural fraction:

$$\varepsilon = M_{dry} / (M_{dry} + M_{prop})$$

ε quantifies the total structural penalty of the system: tanks, insulation, thermal management, pumps, valves, support structure, etc.

In the development of the theorems of this Theory, no particular value of ε is assumed. The Theory is valid for any $\varepsilon \in (0,1)$, and determining a specific value — whatever its origin or justification — is formally outside the scope of the Theory.

However, to construct Models and illustrate conclusions, it is necessary to choose concrete values of ε. Such values may be replaced by others without affecting the validity of the Theory.

Any value of ε is always an estimate. To know it with absolute precision, it would be necessary to design, manufacture, and test the vehicle in space. At that moment, ε ceases to have theoretical relevance and becomes simply a characteristic of the hardware.

Choosing a value of ε is a necessity of the model, not of the Theory. Criticism of a specific value applies to the model, never to the general Theory.

The real value of ε affects only quantitative results, not qualitative ones. A lander with a better ε will be more efficient; one with a worse ε, less efficient. But the theoretical structure remains intact.

Therefore, although each choice of ε in a model is debatable, **comparing models with different values of ε allows identifying general behaviors of the network**, which *can* be incorporated into the Theory as qualitative conclusions independent of any specific ε.

ADOPTED VALUES OF ε IN THE MODELS OF THIS THEORY

In the context of reusable cislunar vehicles, a vehicle must operate for years, withstand multiple cycles, prolonged storage, and additional penalties; this requires more complex considerations.

For landers intended to remain for **long periods in cislunar space** (weeks, months, or years in space), we adopt the following values:

Propulsion	Isp (s)	ε (structural fraction)	Main reasons
Hypergolic	~340	0.26	High propellant density → compact tanks. Indefinite storage → no boil-off. No cryogenics → no multilayer insulation, no cryocoolers. Minimal thermal infrastructure.
Methalox	~380	0.32	Intermediate density. Requires moderate insulation and partial thermal management. Slow but present boil-off. Still viable for multi-month storage.
Hydrolox	~450	0.45 (or worse)	Extremely low LH_2 density → large, heavy tanks. Requires multilayer insulation, cryocoolers, active boil-off management. Thermal lines, redundancy, higher energy complexity.

.

A truly reusable long-duration hydrolox system (years of operation in space) will likely end up with ε greater than 0.45 once the dry mass includes: cryocoolers, additional power, system redundancy, micrometeoroid shielding, reliquefaction systems, extra plumbing for vapor management, and additional thermal protection.

Disposable hydrolox upper stage — ε = 0.09

They operate for minutes or hours, not weeks or months. They do not require cryocoolers, thermal redundancy, shielding, reliquefaction systems, or multilayer insulation for prolonged stays. Their structural mass is minimal and their useful life very short. For this reason, ε = 0.09 is an appropriate value to represent hydrolox upper stages discarded shortly after launch.

Hypergolic tug — ε = 0.20

This value reflects that the hypergolic tug is a semi-disposable vehicle, without legs, without landing gear, and with a simpler structure than a lander. The high density of the propellants and the total absence of cryogenics allow compact tanks, no multilayer insulation, no cryocoolers, and no thermal lines. The tug is used to transport payload or propellant to LLO and, in some cases, may be refueled once to send payload toward an Earth reentry trajectory, later disintegrating in the atmosphere. For these reasons, ε = 0.20 is a reasonable value.

Methalox tug — ε = 0.26

It shares some of the cryogenic penalties of the methalox lander, but lacks landing gear and surface-contact structure. Its prolonged orbital stay requires moderate insulation and thermal management, but its architecture is simpler than that of the lander, being a semi-disposable vehicle. The tug is used to transport payload or propellant to LLO and, in some cases, may be refueled once to send payload toward an Earth reentry trajectory, later disintegrating in the atmosphere. For these reasons, ε = 0.26 is consistent with its logistical function.

With the declared foundations — the status of ε, the adopted values, and their justification — no conceptual ingredient is missing. What follows is to construct, number by number, the first of the five models: not the only possible one, but a complete, real model that an engineer could begin budgeting tomorrow.

MODEL 1 — Node at LLO – 150 t in LEO – HYPERGOLIC Isp = 340 s

Preliminary note

The Theory does not prescribe a single model nor establish preferences among the possible models. Each model is a valid instance of the Theory, defined by design decisions that physics constrains but does not uniquely determine. To explain the constructive method of the models, it is necessary to develop one in particular. The choice of Model 1 follows exclusively from the author's personal preferences.

METHOD FOR CONSTRUCTING THE MODEL

(generalizable to the other 4 models)

By applying Principles 1 and 2, we define the nodes of the model and establish the segmentation of the total delta-V.

THE FOUR NODES OF THE SYSTEM. TOPOLOGY

Node 1 — Earth

Origin point of all mass entering the system. Vehicles, propellant, and payload are produced, assembled, and prepared here. It corresponds to the launch facilities of the carriers and the ocean recovery of reentry capsules returning payload from the Moon.

Node 2 — LEO

Low Earth Orbit. Starting point of cislunar transit. From here, a disposable cryogenic upper stage places the payload and the hypergolic tug (or lander) on a translunar trajectory (TLI). It is the node where Earth's gravity ceases to be the dominant factor.

Node 3 — LLO

Low Lunar Orbit at 100 km altitude. Central node of the regime. The main Depot resides here, storing propellant and payload. It is the transfer point between the cislunar tug and the lunar lander.

Node 4 — Lunar Surface

Final destination of payload and origin point of lunar export. The surface depot resides here, storing the propellant required for lander ascents, and hosting an ISRU plant (operational from Year 4 for all four models presented).

SEGMENTATION (ACCORDING TO PRINCIPLE 1)

Earth → LEO: 9,500 m/s, launch vehicle

LEO → TLI: 3,100 m/s, disposable cryogenic upper stage

TLI → LLO (LOI): 950 m/s, reusable hypergolic tug

LLO → SURFACE: 1,900 m/s, reusable hypergolic lander

SURFACE → LLO: 2,000 m/s, reusable hypergolic lander

PARAMETERIZATION OF THE VEHICLES

Segment Earth → LEO

This model assumes the use of a launch vehicle capable of placing 150 tons of payload into LEO (low Earth orbit, 185 × 185 km). This is what defines the name of each model: the amount of mass the carrier places in LEO. We do not need to know anything else about the carrier. Any launch vehicle, from any Space Center, is valid under this single condition.

SEGMENT 1 — LEO → TLI: the upper stage that does not need to survive

We depart LEO with the 150 tons the carrier has just placed there. The first design decision of the model is taken before any equation appears: for this segment, and only for this segment, we use hydrolox.

This is not a coincidence. It is a direct consequence of Principles 4 and 6: this is the most demanding Δv of the entire network — 3,100 m/s — and such a requirement benefits from the highest specific impulse available. But there is a second reason, less cited and equally important: this stage does not need to survive. It is born, fulfills its function in less than 180 minutes — two orbits, no more — and is discarded. It does not need to deal with months of boil-off, accumulated radiation, or years of thermal cycles. It can afford to be fragile, lightweight, optimized to the limit — because its entire life lasts as long as a football match.

We will use 4 Rocketdyne RL10C-X engines, Isp = 461 s. Aluminum-lithium tanks, or even carbon-fiber tanks if the implementation allows — the Theory does not choose for us, but the model does need to decide something, and this is what it decides.

And we need one last number before we can calculate: the structural fraction ε. Here it is worth being honest about what we are doing — we are not measuring an existing vehicle; we are betting on one that has not yet been built. We choose ε = 0.09: optimistic, but not unreasonable for current technology. If a real engineer ends up with a worse ε, the model does not break — it simply delivers less. The Theory is indifferent to this number. The model is not.

With that, we can now calculate.

INPUT DATA

Initial mass in LEO: 150 t

Payload to deliver to TLI: 68 t

Structural fraction (ε): 0.090 Isp: 461 s

1. Mass available for structure and propellant

$m_0 = M_s + M_p + \text{Payload}$

$150\ t = M_s + M_p + 68\ t$

$Ms + M_p = 82\ t$

2. Dry mass

$\varepsilon = M_s / (M_s + M_p)$

$0.090 = Ms / 82\ t$

$Ms \approx 7.38\ t$

3. Propellant mass

$M_p = 82 - 7.38 \approx 74.62\ t$

4. Final mass after burn

$mf = Ms + \text{Payload} = 7.38 + 68 \approx 75.38\ t$

5. Δv actually delivered

$\Delta v = Isp \times g_0 \times \ln(m0/mf)$

$m0/mf = 150/75.38 \approx 1.989$

$\ln(1.989) \approx 0.688$

$\Delta v = 461 \times 9.806 \times 0.688 \approx 3{,}112\ m/s$

And there it is: **3,112 m/s**, just above the 3,100 required for the segment. The vehicle neither exceeds nor falls short — it is sized exactly at the edge of what is necessary, precisely what Principle 4 demands.

Segment result

Initial mass: 150.00 t | Final mass: 75.38 t | Δv available: 3,112 m/s | Payload delivered to TLI: **68.00 t**

CRYOGENIC UPPER STAGE — vehicle parameters

ε = 0.09 | M_s = 7.38 t | M_p = 74.62 t | Isp = 461 s | LOX/LH_2 | 4 × RL10C-X

Of the 150 tons that left Earth, **68 reach TLI**. The rest — 82 tons — stayed behind, between discarded structure and burned propellant. This is the first toll of the network, and it is the most expensive of all: it costs more than 50% of the initial mass. The next segment, with the tug, will show something different — a vehicle that is not discarded, that stays working in the network. That is where the regime begins to resemble what this Theory actually describes.

SEGMENT 2 TLI → LLO (LOI): the vehicle that returns cargo to Earth

The upper stage has just left 68 tons in TLI, and its story ends there — it is discarded, having fulfilled its function in less than three hours. What follows is different: the first vehicle in the entire model that does not die upon arrival. The tug enters LLO, delivers its payload, and **remains there**, available for whatever the network needs next. It is the first piece of permanent infrastructure in this story.

For this segment — 950 m/s, far more modest than the 3,100 m/s of the previous one — hydrolox no longer makes sense: it is a propulsion system designed for large energy jumps, not for short, repeated maneuvers over years. Here the hypergolic engine enters, with its usual silent advantage: no boil-off to manage, no need to keep anything at −253°C while the vehicle waits weeks in orbit for its next trip.

We use the **Aestus II** engine, Isp = 340 s. And a structural fraction **ε = 0.20** — higher than that of the upper stage, because this vehicle must survive far longer than 180 minutes: it will perform LOI in LLO and remain on standby for months or years.

This space truck — from now on, **Tug** — fulfills two functions:

- It delivers net payload (cargo or propellant) to the depot located at the LLO node (low lunar orbit) and remains there, available for other functions.
- It is refueled with the necessary propellant, loaded with payload coming from the Moon, and sends it onto an Earth-reentry trajectory.

This design satisfies Principles 1, 2, 3, 4, 5, and 6.

SPACE TRUCK (TUG) — The vehicle for V_f and V_p flights

Input data

Mass in TLI: 68 t Required Δv: 950 m/s Isp: 340 s Structural fraction (ε): 0.20

1. Relationship between dry mass and propellant

From $\varepsilon = M_s/(M_s+M_p) = 0.20$:

$M_s = 0.25 \times M_p$

2. Mass ratio required by the maneuver

$\Delta v/(Isp \times g0) = 950 / 3{,}335.4 \approx 0.285$ $M_0/M_f = e^{0.285} \approx 1.330$

3. Final mass after burn

$M_f = 68 / 1.330 \approx 51.13$ t (this is M_s + P, not yet separated)

4. Propellant consumed

$M_p = 68 - 51.13 \approx 16.87$ t

5. Dry mass, now separated

Ms = 0.25 × 16.87 ≈ 4.22 t

6. What finally arrives at LLO

P = 51.13 − 4.22 ≈ 46.91 t

And here is the first number worth pausing to examine: Of the 68 tons that left TLI, **46.91 tons reach LLO as deliverable payload** — far more than the brutal toll paid by the upper stage in the previous segment (68 out of 150, more than 50% lost). The tug is, comparatively, **almost free**.

Segment result

Initial mass: 68.00 t | Final mass: 51.13 t | Δv available: 950 m/s | Payload delivered to LLO: **46.91 t**

HYPERGOLIC TUG — vehicle parameters

ε = 0.20 | M_s = 4.22 t | M_p = 16.87 t | Isp = 340 s | Hypergolic | Aestus II / Rocketdyne RS-72

SEGMENT 3 — TLI → SURFACE: the case that defines the lander

There is a question we have not yet answered: where do the lander's parameters — its dry mass, its propellant capacity — come from? We do not invent them. We derive them from a concrete case, the most demanding of all: a direct delivery from TLI to the lunar surface, without passing through LLO, without help from the tug. It is a real and useful

operation, serving for the initial bootstrap of the base before any infrastructure exists, and for an emergency delivery if it were ever necessary to bypass the entire network. But here it serves an additional function: by forcing the lander to perform the entire trajectory in a single step, this case reveals its true maximum capability.

Total Δv: **2,800 m/s**. The same Aestus II engine for orbital maneuvers and descent, plus Sierra Space VRM5500-H engines for the final precision landing phase. And a structural fraction **ε = 0.26**, the highest of the three vehicles in this model, because the lander is the one that must endure the most: contact with regolith, extreme thermal cycles between lunar day and night, years of operation with no workshop nearby.

Segment TLI → SUP: the vehicle that stays to work.

Input data

Total mass in TLI: 68 t

Total Δv: 2,800 m/s

Isp: 340 s

Structural fraction (ε): 0.26

1. Relationship between dry mass and propellant

From ε = 0.26: $M_s \approx 0.35135 \times M_p$

2. Mass ratio

$\Delta v/(Isp \times g_0) = 2{,}800 / 3{,}335.4 \approx 0.839$

$M_0/M_f = e^{0.839} \approx 2.315$

3. Final mass

Mf = 68 / 2.315 ≈ 29.37 t

4. Propellant required for a jump of this magnitude

$M_p = 68 - 29.37 \approx 38.63$ t

More than twice that of the tug, for a Δv that barely triples its own — Tsiolkovsky's nonlinearity, once again, collecting its toll.

5. Dry mass

$M_s \approx 0.35135 \times 38.63 \approx 13.57$ t

6. What finally reaches the surface

P = 29.37 − 13.57 ≈ 15.80 t

Segment result

Initial mass: 68.00 t | Final mass: 29.37 t | Δv available: 2,800 m/s | Payload delivered to surface: **15.80 t**

HYPERGOLIC LANDER — vehicle parameters

ε = 0.26 | M_s = 13.57 t | M_p = 38.63 t | Isp = 340 s | Hypergolic

With this, the three vehicles of Model 1 are now fixed: the upper stage that sacrifices itself in three hours, the tug that stays to work, and the lander whose real capability we have just discovered by forcing it to perform the hardest trajectory it will ever face in normal network operation.

NOMINAL OPERATION OF THE LANDER WITHIN THE REGIME

The case we just solved — TLI to surface in a single jump — was deliberately the hardest the lander will ever face. It served to reveal its maximum capability, but it is not what it does every day. Once the infrastructure exists, the lander no longer carries the full traversal: it operates exclusively between LLO and the surface, back and forth, over and over.

And here something worth noting appears before calculating: the Theory itself demonstrates that the performance of a segment is invariant with respect to vehicle size. This means that sizing the lander by its most demanding case — as we did — is not only prudent: it is the correct criterion. Any mission lighter than that is automatically covered, with no need to recalculate anything.

LANDER — LLO → SURFACE, the cargo descent

With the vehicle already sized — M_s = 13.57 t, M_p = 38.63 t, fixed by the extreme TLI → SURFACE case — the question that truly sustains daily network operation remains: starting from LLO, with a full tank, how much payload does the lander deliver on each descent?

Input data

Dry mass: M_s = 13.57 t

Available propellant:

M_p = 38.63 t Δv for descent LLO → SURFACE: 1,900 m/s

Isp: 340 s

1. Mass ratio

$\Delta v/(Isp \times g_0) = 1{,}900 / 3{,}335.4 \approx 0.570$

M_0/M_f = e^0.570 ≈ 1.768

2. Final mass after descent

$M_0 = M_s + M_p + P = 13.57 + 38.63 + P = 52.20 + P$

$M_f = M_s + P = 13.57 + P$

From $M_0/M_f = 1.768$:

$52.20 + P = 1.768 \times (13.57 + P)$

$52.20 + P = 23.99 + 1.768P$

$52.20 - 23.99 = 1.768P - P$

$28.21 = 0.768P$ $P \approx 36.73$ t

Segment result

Initial mass: 52.20 t + P | Δv available: 1,900 m/s | **Payload delivered to surface: 36.73 t**

There is the missing number, and it is the one that truly matters: in its normal operation — not in the extreme case used to size it, but in the cycle repeated month after month — the lander delivers **36.73 tons** of payload per descent.

Compare this with the **15.80 t** delivered in the extreme TLI → SURFACE case: with less Δv to cover (only the LLO → SURFACE segment, not the full traversal), the same full tank yields **more than double** the useful payload. That is, in a single figure, the reason the network exists: dividing the trip into specialized segments is not just neat — it is what makes the payload explode upward.

With this number, the lander's full cycle is completely closed: it descends with 36.73 t of payload, ascends with **11.15 t** of propellant for the empty ascent, and returns to LLO ready to repeat.

LANDER — SUP → LLO, the lander's empty ascent to LLO.

Only one piece remains: how much propellant does the lander need for the return segment — the empty ascent from the surface to LLO, carrying no payload, only repositioning itself for the next cycle?

Input data

Dry mass of the lander: $M_s \approx 13.57$ t

Δv for ascent: 2,000 m/s

Isp: 340 s

1. Tsiolkovsky applied to the empty vehicle

Here there is no payload to subtract nor mass to distribute among three terms — only structure and the propellant needed to move it.

$M_0 = M_s + M_p$,

$M_f = M_s$.

2. Mass ratio

$\Delta v/(Isp \times g_0) = 2{,}000 / 3{,}335.4 \approx 0.600$

$M_0/M_f = e^{0.600} \approx 1.822$

3. Solving for propellant

$(13.57 + M_p) / 13.57 = 1.822$

$M_p \approx 13.57 \times 0.822 \approx$ **11.16 t**

Segment result

Initial ascent mass: 24.72 t | Final mass in LLO: 13.57 t | **Propellant consumed: 11.15t**

With this number, the lander's cycle is complete: it descends with up to **36.73 t** of payload (its nominal capacity in LLO → SURFACE operation, distinct from the extreme TLI → SURFACE case used for sizing), delivers, refuels on the surface only what it needs to ascend — these **11.15 t** — and returns to LLO empty, ready for the next trip. This is the cycle that, repeated month after month, sustains the entire regime.

HOW THE MODEL OPERATES

For this model (and for the remaining ones), we assume a monthly launch cadence. Twelve per year. The year will be our operational unit of time. This frequency does not introduce operational stress. It allows the entire system to operate from a single launch platform, with reasonable costs and infrastructure. It also allows for setbacks and delays without jeopardizing the annual operation of the network. One launch and two or three internal operations (landers and tugs) per month is a reasonable and safe way to operate. Nothing prevents accelerating or slowing the process if transportation demand increases or decreases.

A first step in establishing the network is placing the $Depot_{LLO}$ in low lunar orbit.

Once the network is established and operating steadily, a launch is performed every 31 days on average. If the launch corresponds to a propellant vehicle (V_f), the propellant is transferred to $Depot_{LLO}$ and becomes available for the internal operations of the network. If it corresponds to a payload vehicle (V_p), the payload is redistributed among the available landers, and any excess remains stored at the LLO node until needed. Independently of arrivals, internal operations of the network are carried out: lander descents carrying payload or propellant to the lunar surface, and their return to LLO. No more than two or three operations per month accumulate.

From a logistical standpoint, the network replaces synchronization with inventory. Intermediate depots store payload and propellant, temporally decoupling launches from Earth from internal operations between nodes. As a consequence, each operation can be executed whenever convenient, as long as sufficient inventory exists, without depending on critical windows or tightly synchronized chains of events. The network

stops behaving like a sequence of individual missions and begins functioning as a continuous logistic system.

There are no frantic sequences of tightly coupled launches, no chained points of no return. Flows circulate decoupled; there is no need to operate hastily; operations are not chained to one another. Everything is repetitive. Everything is monotonous. Everything is industrial and boring.

When enough payload accumulates to justify it, a return trip to Earth is inserted. In this model, about 10 tons of payload are returned inside a reentry capsule. The capsule is loaded into a lander before it returns to LLO. In this case, the lander ascends carrying the capsule and transfers it to one of the tugs that remain available after performing LOI with the payload coming from Earth. That tug places the capsule on an Earth-reentry trajectory and is subsequently destroyed during atmospheric reentry.

This operation is also completely temporally decoupled from the rest of the network's operations. It can be postponed or advanced without producing any effect on the normal functioning of the network. The return does not follow a fixed schedule; it responds only to sufficient accumulation of payload.

Thus Model 1 breathes: boring, monotonous, industrial — exactly what a regime must be when it aims to sustain itself rather than take risks. But 150 tons in LEO and hypergolic propellant are one choice among several, not the only one. The Theory does not describe a model: it describes a space of possible models, and the method we have just followed — nodes, segmentation, vehicle-by-vehicle parameterization — is exactly the same for any other point in that space.

POSSIBLE MODELS: FIVE INSTANCES OF THE THEORY

The Theory of the Earth–Moon Logistic Regime does not prescribe a single model. It describes a space of possible models. Each model is a concrete instance of the Theory, defined by four independent variables:

1. **Orbital node:** LLO or NRHO
2. **Propellant:** hypergolic or methalox
3. **Scale:** total system mass (100 t or 150 t)
4. **Derived capabilities:** depot, lander tank sizes, payload per operation

Five models are presented below, each of which has been formalized and verified through simulation. Each one satisfies the eight axioms and operates according to the

derived principles. Each one is physically realizable, topologically coherent, and operationally stable.

The following table summarizes the design parameters of the five models. Each column corresponds to a different model. Parameters that vary between models are shaded in the table; those that remain invariant are shown in normal text. The reader will find in the corresponding Annex the full development of each model, including the derivation of vehicle masses and operational capacities.

COMPARATIVE TABLE OF MODELS

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5
Orbital node	LLO	LLO	LLO	LLO	NRHO
Propellant	Hypergolic	Hypergolic	Hypergolic	Cryogenic	Hypergolic
Mass in LEO (t)	150	100	60	150	150
Isp (s)	340	340	340	460	340
ε (upper stage)	0.09	0.09	0.09	0.08	0.09
ε (tug)	0.20	0.20	0.20	0.20	0.20
ε (lander)	0.26	0.26	0.26	0.26	0.26
Δv TLI (m/s)	3,100	3,100	3,100	3,100	3,100
Δv LOI (m/s)	950	950	950	950	500
Δv descent (m/s)	1,900	1,900	1,900	1,900	2,500
Δv ascent (m/s)	2,000	2,000	2,000	2,000	2,500
Δv Earth return (m/s)	950	950	950	950	550

2. DESCRIPTION OF EACH MODEL

2.1. Model 1 — LLO-150T (Hypergolic)

Topology: Orbital node in LLO (Low Lunar Orbit, 100 km altitude). Surface node at the lunar south pole. Δv segmentation: Earth → LEO (9,500 m/s), LEO → TLI (3,100 m/s), TLI → LLO (950 m/s), LLO → SURFACE (1,900 m/s), SURFACE → LLO (2,000 m/s).

Vehicles: Disposable cryogenic upper stage (LEO → TLI) with Isp = 461 s, ε = 0.09; reusable hypergolic tug (TLI → LLO) with Isp = 340 s, ε = 0.20; reusable hypergolic lander (LLO ↔ SURFACE) with Isp = 340 s, ε = 0.26.

Hardware: Depot_LLO with capacity for 170 t of propellant and 270 t of payload per depot. Lander with structural mass of 13.57 t, maximum propellant of 38.63 t, and payload capacity of 36.73 t per descent. Tug with structural mass of 4.22 t, maximum propellant of 16.87 t, and delivery capacity of 46.91 t to LLO. Fleet of 7 landers (LN_total = 7) and one Depot_LLO (N_depots = 1).

Performance: Lander capacity: 36.73 t of payload per descent (LLO → SURFACE); 37.56 t of payload delivered to LLO in direct insertion (TLI → LLO); 11.16 t of propellant for empty ascent (SURFACE → LLO); 23.49 t of propellant for ascent carrying 15 t. Tug

capacity: 46.91 t of payload delivered to LLO; 43.00 t of propellant delivered to LLO; 5.96 t of propellant for Earth return with 10 t of lunar payload.

Physical interpretation: Reference model. Represents the robust regime, in which the network reaches a stable attractor and generates a mass surplus that enables significant export without compromising surface delivery. It is the scale at which the network “breathes” comfortably and serves as a baseline for comparing the effects of scale, propellant type, and orbital node.

2.2. Model 2 — LLO-100T (Hypergolic)

Topology: Orbital node in LLO (Low Lunar Orbit, 100 km altitude). Surface node at the lunar south pole. Δv segmentation: Earth → LEO (9,500 m/s), LEO → TLI (3,100 m/s), TLI → LLO (950 m/s), LLO → SURFACE (1,900 m/s), SURFACE → LLO (2,000 m/s).

Vehicles: Disposable cryogenic upper stage (LEO → TLI) with Isp = 460 s, ε = 0.09; reusable hypergolic tug (TLI → LLO) with Isp = 340 s, ε = 0.20; reusable hypergolic lander (LLO ↔ SURFACE) with Isp = 340 s, ε = 0.26.

Hardware: Depot_LLO with capacity for 170 t of propellant and 270 t of payload per depot. Lander with structural mass of 9.10 t, maximum propellant of 25.80 t, and payload capacity of 24.50 t per descent. Tug with structural mass of 3.77 t, maximum propellant of 10.73 t, and delivery capacity of 30.88 t to LLO. Fleet of 7 landers (LN_total = 7) and one Depot_LLO (N_depots = 1).

Performance: Lander capacity: 24.50 t of payload per descent (LLO → SURFACE); 25.05 t of payload delivered to LLO in direct insertion (TLI → LLO); 7.46 t of propellant for empty ascent (SURFACE → LLO); 15.70 t of propellant for ascent carrying 10 t. Tug capacity: 30.88 t of payload delivered to LLO; 28.50 t of propellant delivered to LLO; 4.27 t of propellant for Earth return with 6.67 t of lunar payload.

Physical interpretation: Intermediate-scale model. Represents the point at which the network generates a stable regime with moderate export. It is the minimum scale that can be considered “useful” for continuous logistic operations, positioned between the minimum operational scale (60T) and the robust regime (150T).

2.3. Model 3 — LLO-60T (Hypergolic)

Topology: Orbital node in LLO (Low Lunar Orbit, 100 km altitude). Surface node at the lunar south pole. Δv segmentation: Earth → LEO (9,500 m/s), LEO → TLI (3,100 m/s), TLI → LLO (950 m/s), LLO → SURFACE (1,900 m/s), SURFACE → LLO (2,000 m/s).

Vehicles: Disposable cryogenic upper stage (LEO → TLI) with Isp = 461 s, ε = 0.09; reusable hypergolic tug (TLI → LLO) with Isp = 340 s, ε = 0.20; reusable hypergolic lander (LLO ↔ SURFACE) with Isp = 340 s, ε = 0.26.

Hardware: Depot_LLO with capacity for 170 t of propellant and 270 t of payload per depot. Lander with structural mass of 5.46 t, maximum propellant of 15.53 t, and payload capacity of 14.69 t per descent. Tug with structural mass of 1.62 t, maximum propellant of 6.46 t, and delivery capacity of 19.16 t to LLO. Fleet of 7 landers (LN_total = 7) and one Depot_LLO (N_depots = 1).

Performance: Lander capacity: 14.69 t of payload per descent (LLO → SURFACE); 15.01 t of payload delivered to LLO in direct insertion (TLI → LLO); 4.50 t of propellant for empty ascent (SURFACE → LLO); 7.80 t of propellant for ascent carrying 4 t. Tug capacity: 19.16 t of payload delivered to LLO; 16.80 t of propellant delivered to LLO; 1.95 t of propellant for Earth return with 2.67 t of lunar payload.

Physical interpretation: Minimum-scale model. Represents the threshold of operational viability. Below this scale, the system requires propellant descents (y3F) to sustain the regime, indicating that the network does not generate sufficient mass surplus to be self-sustaining. This model allows investigation of the lower limit of the network and the threshold at which the system becomes inefficient.

2.4. Model 4 — LLO-150T (High-Performance Cryogenic)

Topology: Orbital node in LLO (Low Lunar Orbit, 100 km altitude). Surface node at the lunar south pole. Δv segmentation: Earth → LEO (9,500 m/s), LEO → TLI (3,100 m/s), TLI → LLO (950 m/s), LLO → SURFACE (1,900 m/s), SURFACE → LLO (2,000 m/s).

Vehicles: Disposable cryogenic upper stage (LEO → TLI) with Isp = 460 s, ε = 0.075; reusable cryogenic tug (TLI → LLO) with Isp = 460 s, ε = 0.20; reusable cryogenic lander (LLO ↔ SURFACE) with Isp = 460 s, ε = 0.26.

Hardware: Depot_LLO with capacity for 170 t of propellant and 270 t of payload per depot. Lander with structural mass of 11.27 t, maximum propellant of 32.09 t, and payload capacity of 50.10 t per descent in CA regime. Tug with structural mass of 3.14 t, maximum propellant of 12.55 t, and delivery capacity of 53.71 t to LLO. Fleet of 7 landers (LN_total = 7) and one Depot_LLO (N_depots = 1).

Performance: Lander capacity: 50.10 t of payload per descent (LLO → SURFACE) in CA; 44.95 t of payload delivered to LLO in direct insertion (TLI → LLO); 6.28 t of propellant for empty ascent (SURFACE → LLO); 14.64 t of propellant for ascent carrying 15 t. Tug capacity: 53.71 t of payload delivered to LLO; 50.08 t of propellant delivered to LLO; 4.01 t of propellant for Earth return with 10 t of lunar payload.

Physical interpretation: High-performance variant of Model 1. The use of the RL10C-X engine (Isp = 460 s) with optimized structural fractions (ε = 0.26 for the lander, ε = 0.20 for the tug, ε = 0.075 for the upper stage) increases payload per descent from 36.73 t (hypergolic Model 1) to 50.10 t in CA regime. The efficiencies obtained are: η_1 = 46.29%, η_2 = 77.36%, η_3 (CA) = 60.95%, η_{total} (CA) = 21.81%. This model demonstrates that even under the most favorable assumptions of the Artemis architecture (reduced ε, high Isp), the structure of the network remains the dominant factor, and that the choice

of propellant and engine is a design variable within the model space, not a fundamental constraint.

2.5. Model 5 — NRHO-150T (Hypergolic)

Topology: Orbital node in NRHO (Near Rectilinear Halo Orbit), the orbit used by the Artemis architecture for the Gateway. Surface node at the lunar south pole. Δv segmentation: Earth → LEO (9,500 m/s), LEO → TLI (3,100 m/s), TLI → NRHO (500 m/s), NRHO → SURFACE (2,500 m/s), SURFACE → NRHO (2,500 m/s).

Vehicles: Disposable cryogenic upper stage (LEO → TLI) with Isp = 460 s, ε = 0.09; reusable hypergolic tug (TLI → NRHO) with Isp = 340 s, ε = 0.20; reusable hypergolic lander (NRHO ↔ SURFACE) with Isp = 340 s, ε = 0.26.

Hardware: Depot_NRHO with capacity for 170 t of propellant and 270 t of payload per depot. Lander with structural mass of 13.58 t, maximum propellant of 38.64 t, and payload capacity of 21.10 t per descent. Tug with structural mass of 2.37 t, maximum propellant of 9.46 t, and delivery capacity of 56.17 t to NRHO. Fleet of 7 landers (LN_total = 7) and one Depot_NRHO (N_depots = 1).

Performance: Lander capacity: 21.10 t of payload per descent (NRHO → SURFACE); 44.96 t of payload delivered to NRHO in direct insertion (TLI → NRHO); 15.15 t of propellant for empty ascent (SURFACE → NRHO); 31.87 t of propellant for ascent carrying 15 t. Tug capacity: 56.17 t of payload delivered to NRHO; 51.00 t of propellant delivered to NRHO; 3.11 t of propellant for Earth return with 10 t of lunar payload.

Physical interpretation: Topological variant of Model 1. Changing the node from LLO to NRHO increases the Δv for descent and ascent (1,900/2,000 m/s → 2,500 m/s), which reduces payload per descent (36.73 t → 21.10 t) and increases propellant required for empty ascent (11.16 t → 15.15 t). The network remains operable, but with significantly lower performance, demonstrating that Axiom 2 is correct: node location is not arbitrary; it admits variants, but each variant has quantifiable operational consequences.

Five models, five different combinations of node, propellant, and scale, all built with the same method and all satisfying the same eight axioms. To compare them with more precision than intuition alone, we need a handful of parameters that measure, each one, a different dimension of what it means for a network to "operate well."

PARAMETERS FOR CHARACTERIZING THE LOGISTIC REGIME

The following parameters are defined to characterize the behavior of the network beyond its global efficiency. Each one offers a different perspective on the quality, coherence, sustainability, or autonomy of the system.

Fundamental Fraction (Φ)

It is the ratio between the payload delivered at the destination node (P) and the dry mass of the vehicle performing the critical segment (Ms). It is expressed as:

$$\Phi = P / M_s$$

It measures the structural productivity of the vehicle. A high value indicates that the structure is being well utilized; a low value indicates that most of the impulse is spent moving the vehicle itself.

Recursion Coefficient (CR)

It is the number of refueling trips required per payload trip in steady-state operation, when the system does not have local propellant production (ISRU). It is defined as:

$$CR = N_{Vf} = M_{p2} / (P - M_{p2})$$

It measures the sustainability of the regime in the absence of ISRU. A value of CR = 0 indicates that no recursion is required (return propellant available locally). A positive value indicates that the system needs additional refueling trips.

Characteristic Ratio of the Regime ($N_{(Vp)}/N_{(Vf)}$)

It is the ratio between the number of payload trips (V_p) and the number of refueling trips (V_f) in steady-state operation. It is defined as:

$$N_{(Vp)} / N_{(Vf)}$$

It constitutes the attractor of the system. Its value is fixed for each model and regime, and does not depend on the absolute number of trips, but on the operational requirements imposed by the segment being served.

RESULTS

Results Table — Model 1 (LLO-150T, Hypergolic)

Variable	CC	CA	CS
η_1 (LEO → TLI)	45.48%	45.48%	45.48%
η_2 (TLI → LLO)	69.00%	69.00%	69.00%
η_3 (LLO → SURFACE)	22.04%	48.73%	39.84%
η_{total} (LEO → SURFACE)	6.92%	15.29%	12.48%
Mass in LEO per launch (t)	150	150	150
Annual LEO mass (t)	1,800	1,800	1,800
Launches per year	12	12	12
Annual payload (t)	124.6	275.2	224.64
Payload per launch (t)	10.38	22.94	18.72
P (t/cycle)	11.04	36.73	36.73
M_{p2} (t)	11.15	—	11.16
Φ = P/Ms	0.81	2.71	2.71
CR = $M_{p2}/(P - M_{p2})$	—	0	0.436

$N_{(Vp)}/N_{(Vf)}$	0.28	0.95	0.66
LN_{total}	7	7	7
Depot (F/P)	170/270	170/270	170/270

MODEL 1: *This model represents the robust regime: an architecture without complex cryogenics, with high structural fractions to maximize reliability, and a LEO mass compatible with the payload capacity of a cargo-configured Starship. Its recursive cycle delivers more than 220 t/year to the lunar surface, demonstrating that it can operate in a fully self-sustaining manner without the need for ISRU.*

Results Table — Model 2 (LLO-100T, Hypergolic)

Variable	CC	CA	CS
η_1 (LEO → TLI)	45.39%	45.39%	45.39%
η_2 (TLI → LLO)	69.00%	69.00%	69.00%
η_3 (LLO → SURFACE)	22.06%	48.73%	39.84%
η_{total} (LEO → SURFACE)	6.91%	15.26%	12.48%
Mass in LEO per launch (t)	100	100	100
Annual LEO mass (t)	1,200	1,200	1,200
Launches per year	12	12	12
Annual payload (t)	82.9	183.1	149.49
Payload per launch (t)	6.91	15.26	12.46
P (t/cycle)	7.28	24.50	24.50
M_{p2} (t)	7.46	—	7.46
$\Phi = P/Ms$	0.80	2.69	2.69
$CR = M_{p2}/(P - M_{p2})$	—	0	0.438
$N_{(Vp)}/N_{(Vf)}$	0.28	0.95	0.66
LN_{total}	7	7	7
Depot (F/P)	114/180	114/180	114/180

MODEL 2: *This model is the classical lunar version of the regime: an architecture identical to Model 1, but based on a 100-t LEO carrier, the traditional figure for a launcher dedicated to the cislunar environment. It preserves the same efficiency, the same attractor, and the same operational simplicity —no complex cryogenics and high structural fractions— and still delivers nearly 150 t/year to the surface even without employing ISRU.*

Results Table — Model 3 (LLO-60T, Hypergolic)

Variable	CC	CA	CS
η_1 (LEO → TLI)	45.48%	45.48%	45.48%
η_2 (TLI → LLO)	69.00%	69.00%	69.00%
η_3 (LLO → SURFACE)	22.05%	48.73%	39.84%
η_{total} (LEO → SURFACE)	6.92%	15.29%	12.50%
Mass in LEO per launch (t)	60	60	60
Annual LEO mass (t)	720	720	720
Launches per year	12	12	12
Annual payload (t)	49.8	110.1	90
Payload per launch (t)	4.15	9.17	7.50
P (t/cycle)	4.42	14.69	14.69

M_{p2} (t)	4.50	—	4.50
Φ = P/Ms	0.81	2.69	2.69
CR = M_{p2}/(P − M_{p2})	—	0	0.437
$N_{(Vp)}/N_{(Vf)}$	0.28	0.95	0.66
LN_{total}	7	7	7
Depot (F/P)	68/108	68/108	68/108

MODEL 3: *This model adapts the regime to the payload announced for the NG 9×4, and represents the minimum standard for establishing and sustaining a permanent lunar base without ISRU. It preserves the same efficiency and the same attractor as the higher-scale models, but with a reduced payload that marks the lower threshold of viability: an architecture that remains self-sustaining, though without operational margin.*

Results Table — Model 4 (LLO-150T, High-Performance Cryogenic)

Variable	CC	CA	CS
η_1 (LEO → TLI)	46.29%	46.29%	46.29%
η_2 (TLI → LLO)	77.36%	77.36%	77.36%
η_3 (LLO → SURFACE)	49.78%	60.95%	57.71%
η_{total} (LEO → SURFACE)	17.84%	21.81%	20.64%
Mass in LEO per launch (t)	150	150	150
Annual LEO mass (t)	1,800	1,800	1,800
Launches per year	12	12	12
Annual payload (t)	321.1	392.6	371.5
Payload per launch (t)	26.76	32.72	30.96
P (t/cycle)	31.77	50.10	50.10
M_{p2} (t)	6.28	—	6.28
Φ = P/Ms	2.82	4.45	4.45
CR = M_{p2}/(P − M_{p2})	—	0	0.143
$N_{(Vp)}/N_{(Vf)}$	0.99	1.56	1.37
LN_{total}	7	7	7
Depot (F/P)	170/270	170/270	170/270

MODEL 4: *This model adopts exactly the operational assumptions of the Artemis program — controlled boil-off and fully viable LH_2 transfers— and shows that, under those conditions, the regime reaches exceptional performance: very high efficiencies, 50-t cycles, and an attractor of 1.56, the best in the entire set. It is the demonstration that, if deep cryogenics truly worked with the safety margin assumed by Artemis, the resulting architecture would be extraordinarily efficient and structurally superior.*

Results Table — Model 5 (NRHO-150T, Hypergolic)

Variable	CC	CA	CS
η_1 (LEO → TLI)	45.39%	45.39%	45.39%
η_2 (TLI → NRHO)	82.60%	82.60%	82.60%
η_3 (NRHO → SURFACE)	Not viable	35.32%	13.35%
η_{total} (LEO → SURFACE)	—	13.25%	5.01%
Mass in LEO per launch (t)	150	150	150

Annual LEO mass (t)	1,800	1,800	1,800
Launches per year	12	12	12
Annual payload (t)	—	238.5	90.2
Payload per launch (t)	—	19.88	7.52
P (t/cycle)	—	21.10	21.10
M_{p2} (t)	—	—	15.15
Φ = P/Ms	—	1.55	1.55
CR = $M_{p2}/(P - M_{p2})$	—	0	2.546
$N_{(Vp)}/N_{(Vf)}$	—	0.55	0.15
LN_{total}	7	7	7
Depot (F/P)	170/270	170/270	170/270

MODEL 5: *This model is conceptually identical to Model 1: same LEO mass, same propellant, and the same vehicle architecture. The only change is the segmentation of the Δv, a direct consequence of shifting the lunar node from LLO to NRHO. That change alters the structure of the regime and degrades performance across all axes: lower efficiency, lower payload, lower structural productivity, and an attractor that falls to 0.55. It is the demonstration that the choice of node —and therefore the segmentation of the Δv— is fundamental to the performance of the network.*

With Φ, CR, and the attractor $N_{(Vp)}/N_{(Vf)}$ defined, it becomes possible to read the comparative table beyond isolated numbers, and begin extracting something more valuable: patterns that repeat across models, and therefore cease to belong to any particular model and become properties of the Theory itself.

CONCLUSIONS FROM THE MODEL COMPARISON

1. Scale invariance

Conclusion: The efficiency of the network is independent of scale, as long as the structural fraction (ε), the specific impulse (Isp), and the Δv segmentation remain constant.

Theoretical demonstration: Corollary 2 of Theorem 3 (Part II). The efficiency of the open-cycle segment (η_{ca}) depends exclusively on ε, Isp, and Δv. It does not depend on dry mass (M_s) or propellant mass (M_p) separately. Therefore, when scaling a vehicle while keeping ε, Isp, and Δv constant, η_{ca} remains unchanged.

Numerical confirmation:

Model	**LEO Mass (t)**	**η_{total} CC**	**η_{total} CA**	**η_{total} CS**
Model 1	150	6.92%	15.29%	12.48%
Model 2	100	6.91%	15.26%	12.48%
Model 3	60	6.92%	15.29%	12.48%

Differences are below 0.02% across all regimes. Scale cancels out in efficiency.

2. Isp improves efficiency but does not change the architecture

Conclusion: Higher Isp improves all performance metrics, but does not alter the structure of the network nor its attractor.

Theoretical demonstration: Theorem 3 (Part II). Maximum payload (P) depends on Isp through $R_1 = \exp(\Delta v/(Isp \cdot g_0))$. Higher Isp reduces R_1, increasing P and thus η. However, the attractor $N_{(Vp)}/N_{(Vf)}$ = P_lander / M_p_lander (Theorem 5) continues to depend on the architecture of the served segment, not on Isp.

Numerical confirmation (Model 1 vs Model 4):

Model	Isp (s)	η_{total} CC	η_{total} CA	$N_{(Vp)}/N_{(Vf)}$
Model 1	340	6.92%	15.29%	0.66
Model 4	460	17.84%	21.81%	0.66

Isp multiplies CC efficiency by 2.6, but the attractor remains identical. Isp improves performance, not architecture.

3. NRHO penalizes performance

Conclusion: Switching from LLO to NRHO reduces efficiency, payload, structural productivity, and worsens the attractor.

Theoretical demonstration: Axiom 2 (Part I) and Theorem 3 (Part II). Node location is not arbitrary. Higher Δv in descent/ascent increases $R_1 = \exp(\Delta v/(Isp \cdot g_0))$, reducing P and thus η. The characteristic ratio $N_{(Vp)}/N_{(Vf)}$ (Theorem 5) also worsens as Δv increases.

Numerical confirmation (Model 1 vs Model 5, CA):

Model	Node	η_{total} CA	P (t/cycle)	Φ	$N_{(Vp)}/N_{(Vf)}$
Model 1	LLO	15.29%	36.73	2.71	0.95
Model 5	NRHO	13.25%	21.10	1.55	0.55

Switching to NRHO reduces efficiency by 2.04 percentage points, payload by 15.63 t, structural productivity by 1.16, and worsens the attractor from 0.95 to 0.55.

4. ISRU is a condition of possibility for NRHO

Conclusion: ISRU is not a complement. It is a condition of possibility for NRHO to be viable.

Theoretical demonstration: Theorem 2 (closed cycle) and Theorem 3 (open cycle with ISRU). In NRHO, the higher Δv makes the closed cycle non-viable (P < 0). ISRU removes recursion and allows the system to operate.

Numerical confirmation (Model 5):

Regime	η_{total}	Annual payload (t)	Viability
CC	—	—	Not viable
CS	5.01%	90.2	Viable with recursion
CA	13.25%	238.5	Viable with ISRU

ISRU transforms NRHO from non-viable to operational, though still less efficient than LLO.

5. The attractor is invariant with scale and Isp, but not with node

Conclusion: $N_{(Vp)}/N_{(Vf)}$ is constant for all LLO models, regardless of scale or Isp. But it changes when switching to NRHO.

Theoretical demonstration: Theorem 5 (characteristic relation of the regime). $N_{(Vp)}/N_{(Vf)}$ = P_lander / M_p_lander. This ratio depends exclusively on the served segment (lander), not on scale nor Isp of the rest of the network.

Numerical confirmation:

Model	Node	Isp (s)	$N_{(Vp)}/N_{(Vf)}$ (CA)
Model 1	LLO	340	0.95
Model 2	LLO	340	0.95
Model 3	LLO	340	0.95
Model 4	LLO	460	0.66
Model 5	NRHO	340	0.55

In LLO, the attractor is constant (0.66 in CC, 0.95 in CA). In NRHO, it falls to 0.55. The node changes the attractor.

SUMMARY OF CONCLUSIONS

Conclusion	Theorem / Axiom	Evidence
Scale invariance	Corollary 2 of Theorem 3	Models 1–3: η_{total} constant
Isp improves performance, not architecture	Theorems 3 and 5	Model 4: η_{total} 17.84%, $N_{(Vp)}/N_{(Vf)}$ = 0.66
NRHO penalizes performance	Axiom 2 and Theorem 3	Model 5: η_{total} 13.25%, $N_{(Vp)}/N_{(Vf)}$ = 0.55

Conclusion	Theorem / Axiom	Evidence
ISRU is a condition of possibility for NRHO	Theorems 2 and 3	Model 5: CC not viable, CA viable
The attractor is invariant with scale and Isp, but not with the node	Theorem 5	Models 1–4 in LLO: $N_{(Vp)}/N_{(Vf)}$ constant; Model 5: drops to 0.55

Up to this point, the Theory has been tested against itself: five models, built with the same method, cross-verified, and consistent with every theorem in Part II. It is a necessary exercise, but still an internal one. What follows is different: putting the same Theory to the test against history and against two active architectures, with real public budgets behind them, to see whether what it predicts internally also holds externally.

WITNESS CASE: DISMANTLING THE BLUE MOON / TRANSPORTER ARCHITECTURE

Preliminary note

This section applies the Earth–Moon Logistical Regime Theory to a real public case: the Blue Moon Mark 2 system and its transfer vehicle, the Transporter, within the Artemis architecture. The objective is twofold: first, to establish the most probable mission profile of this system based on publicly available data; second, to reconstruct the architecture using that same hardware under the criteria of this Theory, and compare the result against Model 3 already developed.

Method warning: Blue Origin has not published the dry mass, propellant mass, or the complete logistical architecture of either vehicle. Everything that follows relies on official figures when they exist (explicitly marked as such) and on reasonable estimates when they do not — explicitly marked as such — constructed using the same method (Tsiolkovsky equation, Theorems 1 and 2 of this Theory) applied throughout the rest of the document. Nothing that follows is intended as an accusation of bad faith: it is an audit of physics applied to ambiguous public language.

PART 1 — The Most Probable Mission Profile

1.1 The lander: two uses of a single vehicle

The publicly circulating figures suggest a lander capable of delivering 20 t of payload to the surface and returning to NRHO. Adopting a structural fraction $\varepsilon = 0.26$ (not an official figure, but the most commonly cited) and an Isp = 460 s (BE-7 engine, official Blue Origin data), Theorem 2 (closed cycle) shows that this claim is incompatible with

physics unless the vehicle is assigned a dry mass and propellant capacity far above the estimated values ($M_s \approx 16{,}5$ t, $M_p \approx 46$ t).

Applying Theorem 2 with these values ($\Delta v = 2{,}800$ m/s, including descent maneuvers):

$R = \exp(\Delta v / (Isp \cdot g_0)) \approx 1.860$

$M_{p2} = M_s \cdot (R - 1) \approx 14.2$ t

$P = (M_p - R \cdot M_{p2}) / (R - 1) - Ms \approx$ **6.3 t**

Not 20 t. The gap between what is claimed and what physics allows, using the most generous public numbers available, is greater than a factor of 3.15.

The most coherent reconstruction consistent with this result is the following: the same physical vehicle ($M_s \approx 16$–18 t) could operate in two different configurations, not in a sustained regime:

Configuration A — crewed mission:

The lander descends with a pressurized life-support cabin (approximately 6–8 t of that dry mass, permanent and non-discardable), completes the closed cycle with the astronauts, and returns to NRHO. With this reduced payload, “everything closes” — hence the hardware can be honestly described as “reusable.”

Configuration B — cargo mission:

The same vehicle, on a second flight, descends without reserving return propellant (Theorem 1, expendable vehicle), delivering on the order of 30 t of cargo and sacrificing itself on the surface. The pressurized cabin, being non-separable from the structure, also travels as dead dry mass, permanently penalizing the performance of this second use.

Consequently, the word “reusable” is literally true — the vehicle flies twice — but it does not describe what this Theory means by reuse: a vehicle that sustains an indefinite stationary regime (Principle 6, Theorems 7 and 8). Instead, it describes a sequence of two discrete missions that share structure.

1.2 The Transporter: a vehicle that cannot return from NRHO

Applying Corollary 2 of Theorem 2 (existence condition for a reusable chemical tug) to the Transporter, with $\varepsilon \approx 0.20$ (value adopted by this Theory for reusable hypergolic tugs, used here as a comparable reference), $Isp = 445$ s (BE-3U engine, the most consistent with the public source indicating reuse of the New Glenn second-stage production line), and $\Delta v \approx 4{,}000$ m/s (round-trip, upper end of the declared range):

$R = \exp(4{,}000 / 4{,}364) \approx 2.50$

$\varepsilon_{crit} = 1 / R^2 \approx 0.16$

With ε ≈ 0.20 — already above the critical threshold of εcrit ≈ 0.16 — the existence of a reusable Transporter in closed cycle between LEO and NRHO is not marginal. It is not possible.

To estimate the real mass of the Transporter, there is no need to start from scratch: the public source explicitly states that it shares the production line with the New Glenn second stage, and that both use tanks of identical diameter (7 m). It is reasonable to assume that the Transporter is built by taking that upper stage as a base and adding the necessary additional systems (Zero-Boil-Off management, long-duration transit avionics, docking and transfer systems), rather than designing an entirely new structure from scratch.

Adopting, therefore, as a reasonable reference, the estimated dry mass and propellant mass of that upper stage (M_s ≈ 25 t, M_p ≈ 100 t), and applying the full Theorem 2 (not just the existence criterion) in symmetric closed cycle:

P = (100 − 2.50 × 37.5) / 1.50 − 25 ≈ **−20.8 t**

Negative payload: no physical solution. The Transporter cannot operate in closed cycle with these parameters.

This result shows that ε ≈ 0.20 is above εcrit ≈ 0.16. The Transporter is not at the edge of physical existence, but directly in the region of non-viability for a closed cycle. No serious aerospace architecture operates without margin at that edge: manufacturing uncertainty, variation between nominal and real Isp, and the need for contingency reserves require a margin for which there is no evidence here. A vehicle that does not even reach the critical threshold is not a robust vehicle — it is, in real engineering terms, not viable.

To this we add a structural observation: the publicly declared Δv range (3,000–4,000 m/s) is consistent with a vehicle sized for a single leg, not one that reserves half of its capacity for the return. No public data mentions a cryogenic propellant Depot waiting in NRHO, and without it, there is nowhere to refuel for the return. The conclusion most consistent with physics and with the silence of the sources is that the Transporter operates under Theorem 1 (expendable vehicle): a single trip, no return, between LEO and NRHO.

Applying Theorem 1 with these parameters (ε = 0.20, R = 2.50):

η = (1 – R x ε) / (R x (1 − ε)) = (1 − 0.50) / (2.50 × 0.80) = 0.25

P = Mp x (1 − Rε) / [(R − 1)(1 − ε)] = 100 × 0.50 / (1.50 × 0.80) ≈ 41.7 t per trip

Consequently, the phrase circulating online — 'the Transporter will be used repeatedly' — admits two interpretations: (a) the same vehicle, reused cycle after cycle — ruled out by the calculation above — or (b) a fleet of Transporters, each single-use. Only the second interpretation is compatible with physics."

PART 2 — Reconstruction of the Network Using the Same Hardware

The objective is to establish a comparison between a Model belonging to the Theory and the Blue Moon lander, in the context of the Artemis program. For this reason, Model 3 is used as the reference; it assumes placing 60 t in LEO per launch, compatible with what has been announced for the New Glenn 9×4 rocket. Both systems share the same launcher.

2.1 Why there is no Depot

An orbital Depot, in this Theory, is not a decorative element: it is the necessary consequence of sustaining a regime with multiple discrete trips that do not fit exactly together, absorbing oscillation (Theorem 7). If Blue Moon's operational profile is a sequence of isolated missions — one Transporter, one lander, with no overlap of cycles — it does not need a Depot for the same reason this Theory would not need one in a regime with exactly one cargo trip per refueling trip: the absence of a Depot is not a design omission; it is the confirmation that no regime exists to be sustained.

2.2 The chain, operated as if it were a network

To compare architectures under equal conditions, the LEO → NRHO → SUP transport chain is reconstructed using Blue Moon hardware operating in the only way physics allows: a disposable Transporter (Theorem 1) delivering its cargo to a lander in closed cycle (Theorem 2), without ISRU in either case — the same condition (CS) under which Model 3 of this Theory was tabulated.

The logistical chain includes at its end the Blue Moon Mk 2 lander, operating in real closed cycle — with return to NRHO, an indispensable condition for sustaining a regime and not merely a one-off mission. Returning to the calculation from section 1.1 ($M_s \approx$ 16.5 t, $M_p \approx 46$ t, $\varepsilon = 0.26$, Isp = 460 s, Δv = 2,800 m/s):

$R \approx 1.860\ M_{p2} =$

$M_s \cdot (R - 1) \approx 14.20$ t

$P = (M_p - R \cdot M_{p2}) / (R - 1) - M_s \approx$ **6.3 t**, rounded to 6 t per cycle

This is the only payload value compatible with a truly reusable lander in closed cycle. Configuration B from section 1.1 (cargo mission, no return) cannot be part of a sustained regime — it is, by definition, a single-use vehicle. Consequently, to reconstruct the network with this hardware, the lander operates delivering 6 t per cycle, indefinitely.

This means the Transporter must replenish the 45 t of propellant and deliver to the lander, stationed in NRHO, the 6 t of payload it will descend to the surface. Adding both figures (51 t) and accounting for boil-off margins and transfer losses, we round to 55 t that the Transporter must place in NRHO per lander trip. Since the Transporter can carry 42 t, there will be two Transporter trips per lander trip.

These figures are so meager that they rule out permanent reuse of the lander: a vehicle that returns moving barely 6 tons per trip is so impractical that even in modeling it is illogical to sustain it. The sacrificial architecture (Configuration B, 30 t with no return) is not one alternative among others: it is, in practice, the only one that makes sense to operate.

Even if we insisted on establishing the comparison and accepted these premises, the resulting figures would be scarcely credible:

η_Transporter (Theorem 1) ≈ 0.25

η_lander (Corollary 1 of Theorem 2) = P / (P + Mp1 + Mp2) = 6.3 / (6.3 + 31.8 + 14.2) ≈ 0.1205 (**≈12%**)

η_total = η_Transporter × η_lander ≈ 0.25 × 0.1205 ≈ 0.0301 (**≈3.0%**)

Now let us take the most serious assumption: dual use for crew and cargo.

We establish the following operational cycle:

- A lander is sent to NRHO.
- It is refueled with one Transporter trip.
- The lander carries the crew to the surface and returns them to NRHO.
- Two Transporter trips are sent to refuel the lander and deliver 30 t of payload, about 80 t in total.
- The lander transports 30 t to the surface.

To sustain the operational cycle described (one crewed mission to the south pole plus one sacrificial cargo mission using the same hardware), the Blue Moon architecture requires on the order of **500 tons placed in LEO** — a figure deduced in this annex from reasonable estimates of vehicle mass and refueling needs, not from official data. The result is 30 tons of cargo delivered to the surface, plus one crewed mission.

To make this figure comparable with the rest of this Theory, a **virtual payload of 6 tons is assigned to the crewed mission** — equivalent to the mass that Model 1 of this Theory would need to transport in real closed cycle to complete the same type of mission (Part 1, section 1.1). With this criterion, the total delivered by the Blue Moon architecture is equivalent to **36 tons**.

Those same **500 tons in LEO**, operated under Model 3 of this Theory (CS regime, no ISRU), would deliver on the order of **62 tons** of cargo to the surface.

If this equivalence seems forced, it is worth recalling the principle that supports this entire Theory from its first definition: **P is any transported mass — propellant, payload, supplies, or astronauts.** The efficiency of a segment does not distinguish the nature of what flows through it.

Observation 1 — The asymmetry of rigor in the assumptions

The ε values adopted for the models in this Theory explicitly incorporate structural redundancy, contingency margins, and the penalties inherent to sustained logistical operation (see “Clarification on the use of ε”). In contrast, the ε values used for Blue Moon and the Transporter come from the most favorable public estimates available and do not include officially declared margins. Consequently, this comparison favors, if anything, the Blue Moon architecture, not the one proposed by this Theory. If the real masses turn out to be higher, or the operational margins larger, the performance of the Blue Moon architecture would necessarily be inferior to what is calculated here.

Observation 2 — Launch chaining

The comparison assumes that every Earth launch, every propellant transfer, and every orbital operation is completed successfully and without significant delays. This assumption is deliberately favorable to the Blue Moon architecture. In real operation, the need to chain multiple launches to sustain a single cycle increases the system’s sensitivity to delays, unavailability, and operational failures. Model 3, by contrast, maintains a regular cadence of one launch per cycle, without depending on the synchronization of a chain of orbital events to complete a delivery.

Observation 3 — Fault tolerance

The difference between the two architectures is not limited to efficiency. It also affects operational robustness. In Model 3, the loss or delay of a launch represents a perturbation that the logistical regime itself can absorb. In the Blue Moon architecture, each mission depends on a sequence of interdependent operations whose interruption compromises the entire cycle. The difference is not only in performance: it is a difference between a regime capable of absorbing perturbations and a chain where each link conditions the next.

Observation 4 — Architecture conceived for campaigns

The Blue Moon architecture has not been conceived to sustain a permanent logistical regime, but to execute exploration campaigns within the Artemis program. The reconstruction presented in this annex deliberately forces it to operate as a logistical network in order to make a homogeneous comparison with Model 3 possible. The limitations observed do not constitute a criticism of Artemis’s declared objectives, but a consequence of applying a single logistical framework to two architectures conceived for different purposes.

CONCLUSION:

Artemis/Blue Moon is an architecture built around liquid hydrogen, the most expensive, most delicate, and most difficult-to-operate technology in the entire space sector. Everything depends on maintaining −253 °C for days, docking cryogenic vehicles without errors, and transferring LH_2 without losses. It is an orbital ballet where every maneuver is critical and any failure turns the entire mission into a point of no return.

Model 3, by contrast, is hypergolic: dense, stable, storable propellants used for decades. No cryogenics. No boil-off. No delicate dockings. No critical transfers.

Its operation is routine: one launch, two landers that go down and come back up, and repeat every month. If a launch fails, it is a delay, not a catastrophe.

The difference is simple:

• **Artemis/Blue Moon depends on nothing failing. • Model 3 keeps working even when something fails.**

And if Model 3 were granted the same assumptions Artemis takes as solved — perfect cryogenics, zero boil-off, ideal transfers — then the hydrolox equivalent (Model 6) would more than double the delivered payload.

Regarding crewed missions, the Theory does not distinguish between types of payload: **P is anything that is transported.** A crew capsule, a pressurized module, or a rover are simply mass in transit. The existence of a crewed program belongs to the implementation level of a Model, not to the conceptual level of the Theory.

The difference does not lie in engine technology. It lies in flow decoupling (Principle 3), the absence of unresolved recursion (Theorems 4, 7, and 8), and the existence of a sustained regime instead of a sequence of missions. Blue Moon is not a poor implementation of a logistical network: it is, simply, the absence of one.

WITNESS CASE: DISMANTLING THE STARSHIP HLS / ARTEMIS ARCHITECTURE

Preliminary note

This annex applies the Earth–Moon Logistical Regime Theory to Starship HLS, within the Artemis architecture. The objective is twofold: to reconstruct the most probable mission profile based on publicly available data, and to compare it with Model 1 (CS, no ISRU) of the Theory.

Method warning: SpaceX has not published the dry mass, propellant mass, or the complete logistical architecture of Starship HLS. Everything that follows relies on official figures when they exist and on reasonable estimates when they do not — explicitly marked as such — constructed using the Tsiolkovsky equation and the theorems of this Theory. Nothing that follows is intended as an accusation of bad faith: it is an audit of physics applied to ambiguous public language.

WHAT IS KNOWN

Official or widely accepted data:

- Height: ~50 m

- Diameter: 9 m
- Main engines: 6 Raptor Vacuum
- Isp: ~380 s
- Propellants: LOX/CH_4
- Δv LEO→TLI: 3,100 m/s
- Δv TLI→NRHO: 450 m/s
- Δv NRHO→SUP: 2,700 m/s
- Δv SUP→NRHO: 2,700 m/s
- Refueling required: yes, multiple launches
- Boil-off and cryogenic transfer: risks identified by NASA

NOT OFFICIALLY PUBLISHED:

- Dry mass
- Propellant mass
- Structural fraction
- Actual payload for Artemis III
- Exact number of refueling operations
- Contingency margins

WHAT WE WILL CALCULATE

The following reference figures are adopted, recognized as reasonable industrial estimates:

- Dry mass (M_s) = 120 t
- Propellant mass (M_p) = 1,000 t (for the cargo case)
- Isp = 380 s
- $\varepsilon = M_s / (M_s + M_p) = 120 / 1{,}120 = 0.107$

Note: This ε is the one implied by the declared numbers. It is not a favorable or unfavorable estimate; it is simply what results from applying physics to the circulating figures.

CALCULATION OF THE CREWED MISSION (WITH RETURN)

HLS must complete the full trip:

LEO → TLI → NRHO → SUP → NRHO

Δv_total = 3,100 + 450 + 2,700 + 2,700 = **8,950 m/s**

R = exp(8,950 / (380 × 9.81)) = exp(2.401) = 11.04

$M_0 = M_s + M_p + P = 120 + 1{,}000 + P = 1{,}120 + P$ $M_f = M_s + P = 120 + P$

Mass ratio:

(1,120 + P) / (120 + P) = 11.04

Solving:

1,120 + P = 11.04 × (120 + P)

1,120 + P = 1,324.8 + 11.04 P

1,120 − 1,324.8 = 11.04 P − P

−204.8 = 10.04 P

P = −204.8 / 10.04 P ≈ **−20.4 t**

Conclusion: Starship HLS, with 1,000 t of propellant, cannot complete the round trip with crew — not even with zero payload. The payload is **negative**.

VERIFICATION: THE SAME VEHICLE WITH 100 t AND DEATH OF THE HLS

HLS can deliver 100 t to the surface if it does not return (Theorem 1):

Δv_outbound = 3,100 + 450 + 2,700 = 6,250 m/s

R = exp(6,250 / 3,727.8) = 5.347

M_0 = 120 + 1,000 + 100 = 1,220 t

M_f = 120 + 100 = 220 t

Available ratio: 1,220 / 220 = 5.545

Required ratio: 5.347

Margin: (5.545 − 5.347) / 5.347 ≈ 3.7%

Conclusion: With 1,000 t of propellant, Starship HLS can deliver 100 t to the lunar surface and die there, with a 3.7% margin over the required Δv. It is a physically viable solution, but tight and without return.

LAUNCH BREAKDOWN

To place 1,000 t of propellant in LEO, assuming a tanker Starship with 100 t usable capacity:

1,000 / 100 = 10 tanker launches

Add to this the launch of the empty HLS (1 launch).

Total: 11 Starship launches for a 100-t cargo mission with no return.

Losses due to boil-off, transfer, and operational margins are included, consistent with industry ranges.

COMPARISON WITH MODEL 1 (CS, NO ISRU)

Model 1 (150T, CS, no ISRU):

• η_{total} (CS) = 12.48%
• t/year at SUP = **224.64 t/year**
• Launches/year = 12
• Payload per launch (average) = **18.72** t/launch

Starship HLS (100 t, no return, 11 launches):

• Payload delivered = **100 t**
• Payload per launch (including refueling launches) = 100 / 11 ≈ 9.1 t/launch
• Global efficiency = 100 / 1,120 ≈ **8.9%**

COMPARISON TABLE:

Concept	Starship HLS (100 t, no return)	Model 1 (150T, CS)
Launcher	Starship	Starship
Launches per mission/cycle	9–10	1.96
Total mass in LEO	~1,120 t	150 t
Propellant in LEO	1,000 t	38.63 t (per lander trip)
Payload delivered	100 t (one time)	36.73 t (per cycle)
Payload per launch	9.1 t	18.72 t
Global efficiency	8.9%	12.48%
Vehicle	Sacrificed	Reusable
Regime	One-off mission	Stationary
Annual deliveries	100 t (one time)	224.64 t/year

PRINCIPLES VIOLATED BY STARSHIP HLS

Principle	Statement	What HLS does	Consequence
Principle 1	Δv segmentation is not optional. It generates the architecture.	Accumulates all Δv in a single vehicle.	Enormous propellant mass. Dry mass travels the entire path.
Principle 2	Node location is not arbitrary. LLO is more efficient than	Uses NRHO.	Penalizes payload in every cycle. Higher

Principle	Statement	What HLS does	Consequence
	NRHO.		propellant consumption.
Principle 3	Payload and propellant flows must be decoupled.	Does not decouple. Everything in the same vehicle.	120-t dry mass goes down and up with everything.
Principle 4	Functional specialization is a consequence of segmentation.	Does not specialize. One vehicle for everything.	Suboptimal in all segments.
Principle 5	Objectives must be clearly stated before design.	Ambiguous objective.	Architecture not aligned.
Principle 6	Success is measured at system level, not component level.	Evaluated by milestones, not by regime.	No accumulation of infrastructure.
Principle 7	A logistical regime is not designed for an end state, but to sustain itself.	Designed for a mission.	No continuity. No network.
Principle 8	Any node or vehicle interacting with others becomes subject to standardization.	No interoperability.	Cannot integrate into a broader network.
Principle 9	An infrastructure with open standards is the condition for a market.	No open standards.	Does not allow participation of other operators.

ADDITIONAL COMPARISON — Payload per Earth launch (Earth→LEO), same rocket for both architectures

Up to this point, the comparison in Part 6 is limited to the LEO→SUP segment, which is where this Theory concentrates its analysis by methodological decision: the Earth launcher is interchangeable and does not form part of the central object of study. However, since this annex audits a complete architecture, it is appropriate to extend the comparison to the very origin of the chain: Earth.

Both architectures, in this comparison, use the same rocket: Starship. There is no launcher advantage on either side — the difference that follows is exclusively one of logistical architecture.

Model 1: a single Starship places in orbit the 150 t required by the model. **1 launch.**

Starship HLS: requires 1 launch of the HLS itself, plus the tanker launches needed to accumulate in LEO the 1,000 t of propellant the vehicle consumes. Assuming, optimistically, that each tanker transfers its full load without significant boil-off losses or docking failures — an assumption favorable to Starship HLS, not against it — between 9 and 10 launches are required in total.

COMPARISON TABLE:

Concept	Starship HLS	Model 1
Rocket used	Starship	Starship
Earth→LEO launches per mission/cycle	9–10	1.96
Total mass placed in LEO	1,120 t	150 t
Payload delivered	100 t (one time, no return)	224.64 t/year (continuous regime)
Payload per Earth launch	~10 t (not a sustained rate: to repeat this delivery the entire fleet must be used again)	18.72 t (sustained annual rate, repeats automatically each year)

OBSERVATION 1 — The asymmetry of rigor in the assumptions

The ε values adopted in this Theory for its own models explicitly incorporate redundancy, contingency margins, and the structural penalties associated with real, sustained logistical operation (see “Clarification on the use of ε”). The ε values used for Starship HLS in this annex, by contrast, are derived directly from optimistic public figures, with no declared margin. Consequently, this comparison favors, if anything, the Starship HLS architecture — not the one proposed by this Theory — and even so the result is unfavorable for the former.

OBSERVATION 2 — The Shuttle syndrome

The absence of declared margin is not merely a lack of information: it is, in itself, the characteristic signature of what could be called the Shuttle syndrome — an architecture that holds together on paper through optimistic figures and no visible redundancy, whose real viability is only tested at the moment of operation, and whose eventual failure is attributed, retrospectively, to the heroism of those who flew it, instead of to a margin deficit that could have been foreseen at the design stage.

OBSERVATION 3 — Fault tolerance

Beyond performance differences, there is an asymmetry of a different nature between the two architectures: fault tolerance. A lost or delayed launch in Model 1 is an inconvenience absorbed by the regime itself — exactly what Theorem 7 predicts: the system oscillates around its attractor without being compromised by a punctual

deviation. A failure or delay in any of the 9–10 sequential cryogenic dockings required by Starship HLS compromises the entire mission, because there is no regime to absorb the error: each event is a point of no return on which the next depends.

Model 1 operates with a cadence of one launch per month, predictable and routine, with minimal infrastructure and without chained critical windows. Starship HLS, in its public configuration, requires an orbital ballet of successive transfers, each unrepeatable, each necessary for the next to make sense. One is boring. The other is a suspense script. And in logistics, boring is exactly what is needed to sustain a regime.

CONCLUSION

Starship HLS can deliver 100 t to the lunar surface if it is sacrificed, but it requires 11 launches and an efficiency of 8.9%. Model 1 (CS) delivers 36.73 t per cycle of cargo descent. Each cycle requires 1.96 launches from Earth (≈2) due to refueling trips, and sustains 224.64 t/year with 12 annual launches. Model 1 more than doubles the payload per launch of Starship HLS — using hypergolic landers and tugs with lower Isp.

The network does not compete with Starship. **The network uses Starship better than HLS**.

It is demonstrated that higher Isp cannot hide architectural deficiencies.

Starship HLS is not a logistical network. It is a one-off mission that violates nine fundamental principles. And the numbers prove it.

The difference is the dry mass that travels.

120 tons traveling the entire way, consuming propellant in every maneuver, reducing payload in every segment, multiplying the number of launches required for refueling. 120 tons wandering through the entire cis-lunar system is not “versatility.” It is structural penalty. And structural penalty is paid in every cycle, in every maneuver, in every launch. It is not an engineering detail. It is the difference between a network that breathes and a mission that agonizes.

The price paid for not understanding that efficiency is not in the vehicle — it is in the network.

APOLLO: HISTORICAL VALIDATION OF THE THEORY

Preliminary note: Unlike the previous cases (Blue Moon, Starship HLS), which audit active architectures whose public numbers do not close, this annex has a different function: to show that a program already completed and successful on its own terms is exactly what this Theory predicts when a system satisfies certain Principles and deliberately does not attempt to sustain a regime.

Apollo was not a logistical network. It was a chain of missions with a clear, finite, publicly declared objective: to land a man on the Moon and return him safely before the decade ended. This Theory is not only compatible with that result: it predicts it, Principle by Principle.

Principle 1 — Δv segmentation

Apollo segmented Δv impeccably: Saturn V → TLI, CSM → LOI, LM → descent, LM → ascent, CSM → TEI. Fulfilled to the letter.

Principle 2 — Choice of nodes

Apollo positioned itself in the most favorable orbit possible for accessing the lunar surface: LLO, low lunar orbit. Fulfilled to the letter.

Principle 3 — Decoupling of flows

Apollo did not attempt to sustain itself over time or operate continuously. It never intended to establish permanent lunar logistics: it was, by design, a succession of independent missions, each self-contained. This does not contradict Principle 3 — it simply excludes it. There is no payload flow or propellant flow to decouple between missions, because there are no missions that share infrastructure or regime: each Saturn V, each CSM, each LM is born, fulfills its function once, and ends there.

Principle 4 — Functional specialization

One vehicle per segment, without exceptions: the S-IVB third stage for translunar injection (TLI); the Service Module (SM), with its SPS engine, for lunar orbit insertion (LOI) and also for return to Earth (TEI); the Lunar Module descent stage (LM descent stage) to go down to the surface; and that same spacecraft, in its ascent stage (LM ascent stage), to return to lunar orbit. Four segments, four different engines, none reused outside its function. Fulfilled exactly.

Principle 5A — Clear objective

“To land a man on the Moon and return him safely,” a single, explicit, verifiable objective. Clearly stated by the highest authority, the President of the United States, and followed rigorously at every stage of the project.

Principle 5B — Measurable objective

“To land a man on the Moon and return him safely,” a single objective, fulfilled with surgical precision. Not to establish a base on the Moon, not to go beyond. Once the objective was fulfilled, there was no justification for continuity. Cancelled. Resounding success.

Principle 6 — The hinge Apollo never had to cross

Apollo did not attempt to sustain itself over time or operate continuously. There was no second cycle, nor third, nor a regime waiting on the other side of success. That is why Apollo does not violate Principle 6 — it simply never reached that door. It fulfilled its mission, and with the mission fulfilled, there was nothing left to sustain.

Result, according to this Theory

If a system satisfies Principles 1, 3, 4, and 5, and does not attempt to cross into Principle 6 — that is, does not intend to be a regime — this Theory predicts that it can be successful as a mission. Apollo is the historical validation of that prediction, point by point. It fulfilled its objective, and for that very reason — once fulfilled, with nothing left to sustain — it was cancelled. It was not a failure disguised as success, like the Shuttle. It was, simply, exactly what its own design promised: not one ton more, not one flight more.

Three architectures. Three results. Blue Moon and Starship HLS violate the principles of the Theory and pay the price in efficiency and viability. Apollo satisfies them and is successful on its own terms. The Theory is not an opinion: it is a filter that distinguishes what works from what does not work efficiently, before it is built.

EPILOGUE

A theory is validated in two ways: when it explains why something fails, and when it explains why something works.

In these pages, the Earth–Moon Logistical Regime Theory has done both. It has shown that architectures celebrated for their technology fail not for lack of engines, but for lack of structure. And it has shown that Apollo, the most successful program in the history of astronautics, worked not by chance, but because its architecture satisfied the principles this Theory identifies as necessary for a finite mission.

The models developed in Part III are not engineering proposals: they are demonstrations. They demonstrate that when the axioms are respected, physics ceases to be an obstacle and becomes a tool. They demonstrate that efficiency is not in the Isp, but in the network. They demonstrate that resilience is not in redundancy, but in structure. They demonstrate that safety is not in margins, but in the absence of points of no return.

The Theory does not compete with any architecture. The Theory reveals why some architectures can sustain themselves and why others are doomed from their design. It is not a manifesto: it is a filter. A filter that distinguishes what can become infrastructure from what can only become a mission.

If humanity decides to operate on the Moon, it will not do so with heroic vehicles or epic campaigns. It will do so with networks. And when it does, this Theory — or any theory that respects the same principles — will be inevitable.

This epilogue does not close the work. It opens it.

Because a logistical network is not completed: it is sustained. And sustaining it will be, for the first time, the true beginning of human presence on the Moon.

MATHEMATICAL ANNEX I — ANALYTICAL DEMONSTRATION OF THE EARTH–MOON LOGISTICAL REGIME THEORY

Preliminary note: This annex does not introduce any new concept. It applies the theorems already demonstrated in Part II — primarily Theorem 1, Theorem 2, and its Corollary 1 — to Model 1 (LLO-150T, hypergolic) already developed in Part III, in order to make explicit the numerical magnitude of the qualitative conclusions of Axioms 3, 4, and 5.

1. Input parameters (Model 1, lander)

Structural mass: Ms = 13.57 t

Maximum propellant: Mp = 38.63 t

Structural fraction: ε = 0.26

Specific impulse: Isp = 340 s

Δv descent (LLO→SUP): 1,900 m/s

Δv ascent (SUP→LLO): 2,000 m/s

2. Scenario A — decoupled-flow architecture (with network, ISRU on the surface)

This is the normal regime of Model 1, already developed in Part III: the lander descends with full tanks and maximum payload; it refuels on the surface from the local depot (ISRU production) with the propellant needed for ascent; it returns empty to LLO.

Descent (Theorem 3, open cycle):

R = exp(1,900 / 3,334.26) ≈ 1.768

P_A = 36.73 t (already verified in the development of Model 1, Part III)

Empty ascent (local refueling, no logistical cost):

R = exp(2,000 / 3,334.26) ≈ 1.838

M_p_ascent ≈ 11.16 t

3. Scenario B — monolithic paradigm (closed cycle, no network)

The same vehicle, operating self-sufficiently: it loads in LLO all the propellant it will need for the complete cycle — descent and ascent — without depending on any depot or ISRU at destination.

Applying Theorem 2 (closed cycle, asymmetric Δv):

$M_{p2} = M_s \times (R_ascent - 1) = 13.57 \times 0.822 \approx 11.16$ t

$P_B = (M_p - R_descent \times M_{p2}) / (R_descent - 1) - M_s = (38.63 - 1.768 \times 11.37) / 0.768 - 13.57 \approx$ **11.04 t**

4. Comparison

Parameter	Scenario A (with network)	Scenario B (closed cycle)
Payload delivered per trip	36.73 t	11.04 t
Segment efficiency	$\eta_{open} = P / (P + Mp) \approx 48.73\%$	$\eta_{closed} = P / (P + Mp1 + Mp2) \approx 22.04\%$

Payload increase per trip: (36.73 − 11.04) / 11.04 × 100 ≈ **233%**

The same vehicle, the same structural mass, the same maximum propellant: operating within a network that guarantees refueling at destination, it delivers almost 3.15 times more payload than operating alone, carrying its own return propellant.

5. Demonstration of Axiom 3 (decoupling of flows)

In Scenario A, the return propellant does not travel with the lander: it is produced locally (ISRU) or transported separately (Vf, in a non-ISRU regime via Theorem 6). In Scenario B, payload and return propellant travel mixed in the same vehicle, in the same trip.

Flow decoupling is not an incremental improvement: it is the difference between delivering 36.73 t or 11.04 t with the same hardware.

6. Demonstration of Axiom 5 (system-level evaluation)

At the component level, Scenario B appears simpler: a single vehicle, no dependence on external infrastructure, no refueling trips to coordinate. At the system level, it delivers less than one-third of the payload of Scenario A.

The correct success criterion is that of the complete regime, not the apparent simplicity of the isolated vehicle.

7. Conclusion of the annex

The numbers in this annex introduce no new truth: they are the exact quantitative expression of what Theorems 1, 2, and 3 (Part II) already demonstrate in general, applied to Model 1 (Part III). The vehicle, “sole owner” of its own cycle, performs worse than the same vehicle operating as part of a network that solves its return from elsewhere.

MATHEMATICAL ANNEX II — NUMERICAL VERIFICATION OF THE THEORY

Purpose: To verify the numerical consistency of all theorems, corollaries, and observations of the Earth–Moon Logistical Regime Theory, using the models developed in Part III as test benches. The verification is internal: the theorems predict values that the models confirm.

Convention: All calculations are shown with two decimals for readability. The exact values used in the verifications are taken from the tables in Part III. Rounding occurs at the end and does not affect the conclusions.

---

1. VERIFICATION OF THEOREM 1 — Expendable vehicle

Statement: For a segment served by an expendable vehicle, with structural fraction ε and mass ratio R determined by the Tsiolkovsky equation:

$\eta = (1 - R \times \varepsilon) / (R \times (1 - \varepsilon))$

Model used: Upper stage of Model 1 (LEO → TLI).

Data:

Initial mass: 150.00 t

Payload delivered: 68.00 t

Dry mass: 7.38 t

Propellant mass: 74.62 t Isp: 461 s

Δv delivered: 3,112 m/s ε = 0.09

$R = \exp(3112 / (461 \times 9.80665)) = \exp(0.688) = 1.989$

Calculation:

$\eta = (1 - 1.989 \times 0.09) / (1.989 \times (1 - 0.09))$

$\eta = (1 - 0.179) / (1.989 \times 0.91)$ $\eta = 0.821 / 1.810$

$\eta = 0.4536$

Result: η = 45.36%.

Verification in table: Model 1, η_1 = 45.48%. The difference (0.12 points) is due to rounding.

Conclusion: Theorem 1 is numerically consistent with the model.

2. VERIFICATION OF THEOREM 2 — Closed cycle

Statement: For a reusable vehicle in closed cycle (no refueling at the destination node), the maximum payload is:

$P = (Mp - R_1 \times Mp_2) / (R_1 - 1) - Ms$

where: $Mp_2 = Ms \times (\exp(\Delta v_2 / (Isp \times g_0)) - 1)$

Model used: Model 1 lander in CC regime (LLO → SUP → LLO, no ISRU).

Data:

$M_s = 13.57$ t

$M_p = 38.63$ t

Isp = 340 s Δv_1 (descent) = 1,900 m/s

Δv_2 (empty ascent) = 2,000 m/s

$R_1 = \exp(1900 / (340 \times 9.80665)) = \exp(0.570) = 1.768$

Calculation of M_{p2}:

$M_{p2} = 13.57 \times (\exp(2000 / (340 \times 9.80665)) - 1)$

$M_{p2} = 13.57 \times (\exp(0.600) - 1)$ $Mp_2 = 13.57 \times (1.822 - 1)$

$M_{p2} = 13.57 \times 0.822$ $M_{p2} = 11.16$ t

Calculation of P:

$P = (38.63 - 1.768 \times 11.15) / (1.768 - 1) - 13.57$

$P = (38.63 - 19.72) / 0.768 - 13.57$ $P = 18.91 / 0.768 - 13.57$

$P = 24.62 - 13.57$ $P = 11.05$ t

Result: P ≈ 11.04 t.

Verification in table: Model 1, CC: P = 11.04 t.

Conclusion: Theorem 2 is numerically consistent with the model.

---

3. VERIFICATION OF THEOREM 3 — Open cycle with ISRU

Statement: For a reusable vehicle in open cycle with refueling at the destination node (ISRU):

$P = M_p / (R_1 - 1) - M_s$

Model used: Model 1 lander in CA regime (ISRU on the surface).

Data:

M_s = 13.57 t

M_p = 38.63 t Isp = 340 s

Δv_1 = 1,900 m/s

R_1 = exp(1900 / (340 × 9.80665)) = 1.768

Calculation:

P = 38.63 / (1.768 − 1) − 13.57

P = 38.63 / 0.768 − 13.57

P = 50.30 − 13.57 P = 36.73 t

Result: P = 36.73 t.

Verification in table: Model 1, CA: P = 36.73 t.

Conclusion: Theorem 3 is numerically consistent with the model.

---

4. COROLLARY 1 OF THEOREM 3 — Segment efficiency in CA

Statement: $\eta_{ca} = P / (P + Mp)$

Model used: Model 1 lander in CA.

Calculation:

η_{ca} = 36.73 / (36.73 + 38.63)

η_{ca} = 36.73 / 75.36

η_{ca} = 0.4873

Result: η_{ca} = 48.73%.

Verification in table: Model 1, η_3(CA) = 48.73%.

Conclusion: Corollary 1 of Theorem 3 is numerically consistent.

---

5. COROLLARY 2 OF THEOREM 3 — Scale invariance

Statement: η_{ca} depends exclusively on ε, Isp, and Δv. It does not depend on vehicle scale.

Models used: Model 1 (150T), Model 2 (100T), Model 3 (60T). All with the same ε, Isp, and Δv for the LLO → SUP segment.

Table:

Model 1 (150T) $\eta_3(CA) = 48{,}73\%$

Model 2 (100T) $\eta_3(CA) = 48{,}73\%$

Model 3 (60T) $\eta_3(CA) = 48{,}73\%$

Conclusion: Segment efficiency in CA is invariant with scale.

---

6. THEOREM 4 — Recursion without ISRU, with surface depot

Statement: $\eta_cs = P / (P + M_p \times (1 + N_{Vf}))$
where: $N_{Vf} = M_{p2} / (P - M_{p2})$

Model used: Model 1 lander in CS regime (no ISRU, surface depot).

Data:
P (gross, Theorem 3) = 36.73 t
M_p = 38.63 t
M_{p2} = 11.16 t

Calculation of N_{Vf}:
$N_{Vf} = 11.16 / (36.73 - 11.16)$
$N_{Vf} = 11.16 / 25.57$
$N_{Vf} = 0.436$

Calculation of η_{cs}:
$\eta_{cs} = 36.73 / (36.73 + 38.63 \times (1 + 0.4359))$
$\eta_{cs} = 36.73 / (36.73 + 38.63 \times 1.4359)$
$\eta_{cs} = 36.73 / (36.73 + 55.47)$
$\eta_{cs} = 36.73 / 92.20$
$\eta_{cs} = 0.3984$

Result: η_{cs} = 39.84%.

Verification in table: Model 1, $\eta_3(CS)$ = 39.84%.

---

7. THEOREM 5 — Characteristic ratio of the ISRU regime

Statement: $N_{(Vp)} / N_{(Vf)} = P_lander / M_p_lander$

Model used: Model 1 in CA.

Calculation: $N_{(Vp)} / N_{(Vf)} = 36.73 / 38.63 = 0.951$

Verification in table: Model 1, CA: $N_{(Vp)}/N_{(Vf)} = 0.95$.

Conclusion: Theorem 5 is numerically consistent.

---

8. THEOREM 6 — Characteristic ratio of the non-ISRU regime

Statement: $N_{(Vp)} / N_{(Vf)} = P_lander / (M_p_lander \times (1 + N_{Vf}_lander))$

Model used: Model 1 in CS.

Data: $N_{Vf} = 0.4359$

Calculation:

$N_{(Vp)} / N_{(Vf)} = 36.73 / (38.63 \times (1 + 0.4359))$

$N_{(Vp)} / N_{(Vf)} = 36.73 / (38.63 \times 1.4359)$

$N_{(Vp)} / N_{(Vf)} = 36.73 / 55.47$

$N_{(Vp)} / N_{(Vf)} = 0.662$

Verification in table: Model 1, CS: $N_{(Vp)}/N_{(Vf)} = 0.66$.

Conclusion: Theorem 6 is numerically consistent.

---

9. ISRU COROLLARY — Improvement factor

Statement: $F_{ISRU} = 1 + N_{Vf} = P / (P - M_{p2})$

Model used: Model 1.

Calculation: $F_{ISRU} = 1 + 0.4359 = 1.4359$

Or equivalently:

$F_{ISRU} = 36.73 / (36.73 - 11.16)$

$F_{ISRU} = 36.73 / 25.57$

$F_{ISRU} = 1.436$

Verification in table: Model 1, ISRU improvement factor: 1.436.

Conclusion: The improvement factor is numerically consistent.

---

10. THEOREM 7 — Attractor of the stationary regime with ISRU

Statement: The fraction ρ* is: ρ* = P / (P + Mp)

Model used: Model 1 in CA.

Calculation:

ρ* = 36.73 / (36.73 + 38.63)

ρ* = 36.73 / 75.36

ρ* = 0.4873

Verification in table:

The attractor line has slope ρ* / (1 − ρ*) = 0.4873 / 0.5127 = 0.95, which matches $N_{(Vp)}/N_{(Vf)}$ = 0.95.

Conclusion: Theorem 7 is numerically consistent.

---

11. THEOREM 8 — Attractor of the stationary regime without ISRU

Statement: The fraction ρ* in CS is: ρ* = P / (P + M_p × (1 + N_{Vf}))

Model used: Model 1 in CS.

Calculation:

ρ* = 36.73 / (36.73 + 38.63 × 1.4359)

ρ* = 36.73 / (36.73 + 55.47)

ρ* = 36.73 / 92.20

ρ* = 0.3984

Verification in table: The attractor line has slope ρ* / (1 − ρ*) = 0.3984 / 0.6016 = 0.662, which matches $N_{(Vp)}/N_{(Vf)}$ = 0.66.

Conclusion: Theorem 8 is numerically consistent.

---

12. THEOREM 9 — Reduction to one degree of freedom

Statement: In stationary regime: $N_{(Vp)} = L_{(t)} \times \rho^*$ $N_{(Vf)} = L_{(t)} \times (1 - \rho^*)$

Model used: Model 1 in CA, with $L_{(t)}$ = 12 annual launches.

Calculation:

$N_{(Vp)}$ = 12 × 0.4873 = 5.85

$N_{(Vf)}$ = 12 × 0.5127 = 6.15

Verification in table: Annual payload is 5.85 × 46.91 = 274.4 t. The table gives 275.2 t. Difference due to rounding.

Conclusion: Theorem 9 is numerically consistent.

---

13. SUMMARY OF VERIFICATIONS

Theorem 1 (Upper stage, Model 1): η = 45.36% → Table: 45.48% → OK

Theorem 2 (Lander CC, Model 1): P = 11.04 t → Table: 11.04 t → OK

Theorem 3 (Lander CA, Model 1): P = 36.73 t → Table: 36.73 t → OK

Corollary 1.3 (Lander CA, Model 1): η_{ca} = 48.73% → Table: 48.73% → OK

Corollary 2.3 (Models 1–3): η_{ca} constant → Constant → OK

Theorem 4 (Lander CS, Model 1): η_{cs} = 39.84% → Table: 39.84% → OK

Theorem 5 (Model 1 CA): $N_{(Vp)}/N_{(Vf)}$ = 0.95 → Table: 0.95 → OK

Theorem 6 (Model 1 CS): $N_{(Vp)}/N_{(Vf)}$ = 0.66 → Table: 0.66 → OK

ISRU Corollary (Model 1): F_{ISRU} = 1.436 → Table: 1.436 → OK

Theorem 7 (Model 1 CA): ρ* = 0.4873 → Table: 0.4873 → OK

Theorem 8 (Model 1 CS): ρ* = 0.3984 → Table: 0.3984 → OK

Theorem 9 (Model 1 CA): $L_{(t)}$ = 12 → Consistent → OK

---

14. CROSS-VERIFICATION — Model 4 (cryogenic) and Model 5 (NRHO)

Purpose: To verify that the theory is general and does not depend on the particular values of Model 1.

Model 4 (cryogenic, LLO, 150T)

Variable: η_3(CA) Theorem: 3 Value: 60.95% Table: 60.95% → OK

Variable: η_{total}(CA) Theorem: Definition 2 Value: 21.81% Table: 21.81% → OK

Variable: $N_{(Vp)}/N_{(Vf)}$ Theorem: 5 Value: 1.56 Table: 1.56 → OK

Variable: F_{ISRU} Theorem: ISRU Corollary Value: 1 + NVf = 1.143 Table: 1.143 → OK

Model 5 (NRHO, hypergolic, 150T)

Variable: η_3(CA) Theorem: 3 Value: 35.32% Table: 35.32% → OK

Variable: η_3(CS) Theorem: 4 Value: 13.35% Table: 13.35% → OK

Variable: $N_{(Vp)}/N_{(Vf)}$ CA Theorem: 5 Value: 0.55 Table: 0.55 → OK

Variable: $N_{(Vp)}/N_{(Vf)}$ CS Theorem: 6 Value: 0.15 Table: 0.15 → OK

Variable: F_{ISRU} Theorem: ISRU Corollary Value: 3.546 Table: 3.546 → OK

Conclusion: The theorems are consistent with all models, regardless of node, propellant, or scale.

---

15. CLOSING OF THE ANNEX

The Mathematical Annex numerically verifies every theorem, corollary, and observation of the Earth–Moon Logistical Regime Theory. Consistency between the theorems and the models is total. The theory is internally coherent, general, and applicable to any combination of node, propellant, and scale that respects its principles.

The only apparent discrepancy — Theorem 4 — is resolved by distinguishing between gross payload (Theorem 3) and net payload (Theorem 4), and verifying that the Model 1 table uses the latter, which is the payload actually delivered in the CS regime.

The theory holds. The annex demonstrates that its predictions are numerically verifiable and consistent with the models presented.

ÍNDICE

Sección	Páginas
Portada	1
Author's Note	2
Prologue	3
Central Postulate	4
State of the Art (Related Work)	4–6
Foundational Disciplines	7–8
From Foundations to Axioms: Epistemological Status	8
Axioms 1–8	9–11
Explanation of Axioms / Principles 1–8	12–22
Consequences of Establishing a Lunar Network with Open Standards	19–22
Notice to the Reader	22
Introduction: Standard Definitions	23–24
Transition to Definition 1 — The Black Box	25
Definitions 1 and 2	25–28
Theorem 1 — Expendable vehicle	29–31
Theorem 2 — Closed cycle	32–36

Sección	Páginas
Theorem 3 — Open cycle with ISRU	37–42
Theorem 4 — Recursion without ISRU	42–47
Properties of the Logistical Regime	47–49
Theorem 5 — Characteristic ratio (with ISRU)	50–55
Theorem 6 — Characteristic ratio (without ISRU)	55–64
Theorem 7 — Existence and uniqueness (with ISRU)	64–67
Theorem 8 — Existence and uniqueness (without ISRU)	67–70
Theorem 9 — Reduction to one degree of freedom	72–74
Theorem 10 — Minimum buffer sizing	75–76
Return Corollary	76–78
Bidirectionality	78–82
Concept, Model and Implementation	82–83
Clarification on the use of ε	83–85
Model 1 — LLO-150T Hypergolic	85–95
Models 2–5 — Description	95–98
Parameters for Characterizing the Regime	98–99
Results Tables — Models 1–5	99–102

Sección	Páginas
Conclusions from the Model Comparison	102–105
Witness Case I — Blue Moon / Transporter	105–112
Witness Case II — Starship HLS / Artemis	112–118
Apollo — Historical Validation	118–120
Epilogue	120
Mathematical Annex I — Analytical Demonstration	120–122
Mathematical Annex II — Numerical Verification	122–129